

HP StorageWorks X9000 File Serving Software CLI Reference Guide

This document describes commands provided with X9000 File Serving Software. It is intended for system administrators managing X9300 Network Storage Gateway systems, X9320 Network Storage Systems, and X9720 Network Storage Systems.

HP Part Number: TA768-96032
Published: April 2011
Edition: Fourth



© Copyright 2009, 2011 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgments

Microsoft, Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Revision History

Edition	Date	Software Version	Description
First	November 2009	5.3.1	Initial release of HP StorageWorks X9000 File Serving Software
Second	April 2010	5.4.0	Added or updated the following commands: ibrix_async_replicate, ibrix_auth, ibrix_cifs, ibrix_edquota, ibrix_fm, ibrix_fs, ibrix_fs_ops, ibrix_migrator, ibrix_online_quotacheck, ibrix_replicate, ibrix_supportticket
Third	December 2010	5.5	Added or updated the following commands: ibrix_auth, ibrix_certificate, ibrix_cifs, ibrix_edquota, ibrix_ftpconfig, ibrix_ftpshare, ibrix_httpconfig, ibrix_httpshare, ibrix_httpvhost, ibrix_localgroups, ibrix_localusers, ibrix_migrator, ibrix_ndmpconfig, ibrix_ndmpsession, ibrix_rebalance, ibrix_server, ibrix_tape
Fourth	April 2011	5.6	Updated ibrix_nic, ibrix_snmptrap

Contents

1	General CLI information.....	10
	Executing X9000 Software commands.....	10
	Command syntax.....	10
2	X9000 File Serving Software commands.....	11
	ibrix_activedirectory.....	11
	Description.....	11
	Synopsis.....	11
	Options.....	11
	See also.....	12
	ibrix_at.....	13
	Description.....	13
	ibrix_auth.....	14
	Description.....	14
	Synopsis.....	14
	Options.....	14
	ibrix_autoconnect.....	16
	Description.....	16
	Synopsis.....	16
	Options.....	16
	See also.....	16
	ibrix_certificate.....	17
	Description.....	17
	Synopsis.....	17
	Options.....	17
	See also.....	18
	ibrix_cfrjob.....	19
	Description.....	19
	Synopsis.....	19
	Options.....	20
	ibrix_cifs.....	22
	Description.....	22
	Synopsis.....	22
	Options.....	22
	See also.....	23
	ibrix_cifsconfig.....	24
	Description.....	24
	Synopsis.....	24
	Options.....	24
	See also.....	25
	ibrix_client	26
	Description.....	26
	Synopsis.....	26
	Options.....	27
	See also.....	27
	ibrix_cluster.....	28
	Description.....	28
	Synopsis.....	28
	Options.....	28
	See also.....	28

ibrix_clusterconfig.....	29
Description.....	29
Synopsis.....	29
Options.....	30
ibrix_dbck.....	31
Description.....	31
Synopsis.....	31
Options.....	31
See also.....	31
ibrix_edquota.....	32
Description.....	32
Synopsis.....	32
Options.....	33
See also.....	34
ibrix_event	35
Description.....	35
Synopsis.....	35
Options.....	36
See also.....	37
ibrix_exportcfr.....	38
Description.....	38
Synopsis.....	38
Options.....	38
See also.....	38
ibrix_exportcfrpreference.....	39
Description.....	39
Synopsis.....	39
Options.....	39
ibrix_exportfs	41
Description.....	41
Synopsis.....	41
Options.....	42
See also.....	42
ibrix_fm.....	43
Description.....	43
Synopsis.....	43
Options.....	44
ibrix_fm_tune.....	45
Description.....	45
Synopsis.....	45
Options.....	46
See also.....	47
ibrix_fs.....	48
Description.....	48
Synopsis.....	48
Create a file system.....	48
View information about file systems.....	49
Extend a file system.....	49
Delete a file system.....	50
Migrate segment ownership.....	50
Enable or disable quotas.....	50
Mark a bad segment.....	50
Enable or disable Export Control.....	50
Disable 32-bit mode.....	51
Assign or delete a standby server for the specified segments.....	51

Options.....	51
See also.....	52
ibrix_fs_ops.....	53
Description.....	53
Synopsis.....	53
Options.....	53
See also.....	53
ibrix_fs_tune.....	54
Description.....	54
Synopsis.....	55
Options.....	56
See also.....	57
ibrix_fsck.....	58
Description.....	58
Synopsis.....	58
Options.....	59
See also.....	59
ibrix_ftpconfig.....	60
Description.....	60
Synopsis.....	60
Options.....	60
See also.....	61
ibrix_ftpshare.....	62
Description.....	62
Synopsis.....	62
Options.....	63
See also.....	63
ibrix_haconfig.....	64
Description.....	64
Synopsis.....	64
Options.....	65
ibrix_hba.....	66
Description.....	66
Synopsis.....	66
Options.....	66
ibrix_health.....	68
Description.....	68
Synopsis.....	68
Options.....	69
ibrix_healthconfig.....	70
Description.....	70
Synopsis.....	70
Options.....	70
ibrix_host_tune.....	71
Description.....	71
Synopsis.....	71
Options.....	72
See also.....	73
ibrix_hostgroup.....	74
Description.....	74
Synopsis.....	74
Options.....	75
See also.....	75
ibrix_hostpower	76
Description.....	76

Synopsis.....	76
Options.....	76
See also.....	76
ibrix_httpconfig.....	77
Description.....	77
Synopsis.....	77
Options.....	77
See also.....	78
ibrix_httpshare.....	79
Description.....	79
Synopsis.....	79
Options.....	80
See also.....	80
ibrix_httpvhost.....	81
Description.....	81
Synopsis.....	81
Options.....	81
See also.....	82
ibrix_license	83
Description.....	83
Synopsis.....	83
Options.....	83
ibrix_localgroups.....	84
Description.....	84
Synopsis.....	84
Options.....	84
See also.....	84
ibrix_localusers.....	85
Description.....	85
Synopsis.....	85
Options.....	85
See also.....	86
ibrix_lv.....	87
Description.....	87
Synopsis.....	87
Options.....	87
See also.....	88
ibrix_migrator.....	89
Description.....	89
Synopsis.....	89
Options.....	90
Writing tiering rules.....	90
See also.....	92
ibrix_mount	93
Description.....	93
Synopsis.....	93
Options.....	94
See also.....	94
ibrix_mountpoint	95
Description.....	95
Synopsis.....	95
Options.....	96
See also.....	96
ibrix_ndmpconfig.....	97
Description.....	97

Synopsis.....	97
Options.....	97
ibrix_ndmpsession.....	98
Description.....	98
Synopsis.....	98
Options.....	98
ibrix_nic.....	99
Description.....	99
Synopsis.....	99
Options.....	100
ibrix_online_quotacheck.....	102
Description.....	102
Synopsis.....	102
Options.....	102
ibrix_powersrc	104
Description.....	104
Synopsis.....	104
Options.....	105
See also.....	105
ibrix_profile.....	106
Description.....	106
Synopsis.....	106
Options.....	106
ibrix_pv	107
Description.....	107
Synopsis.....	107
Options.....	107
ibrix_rebalance.....	109
Description.....	109
Synopsis.....	109
Options.....	110
ibrix_server	112
Description.....	112
Synopsis.....	112
Options.....	113
See also.....	114
ibrix_sh.....	115
Description.....	115
Synopsis.....	115
Options.....	115
ibrix_snap.....	116
Description.....	116
Synopsis.....	116
Options.....	117
See also.....	117
ibrix_snap_strategy.....	118
Description.....	118
Synopsis.....	118
Options.....	119
See also.....	119
ibrix_snmpagent.....	120
Description.....	120
Synopsis.....	120
Options.....	121
See also.....	121

ibrix_snmpgroup.....	122
Description.....	122
Synopsis.....	122
Options.....	122
See also.....	123
ibrix_snmptrap.....	124
Description.....	124
Synopsis.....	124
Options.....	125
See also.....	125
ibrix_snmpuser.....	126
Description.....	126
Synopsis.....	126
Options.....	126
See also.....	127
ibrix_snmpview.....	128
Description.....	128
Synopsis.....	128
Options.....	128
See also.....	128
ibrix_stats.....	129
Description.....	129
Synopsis.....	129
Options.....	129
ibrix_supportticket.....	130
Description.....	130
Synopsis.....	130
Options.....	130
ibrix_tape.....	131
Description.....	131
Synopsis.....	131
Options.....	131
ibrix_task.....	132
Description.....	132
Synopsis.....	132
Options.....	132
ibrix_tier.....	133
Description.....	133
Synopsis.....	133
Options.....	133
See also.....	133
ibrix_umount.....	134
Description.....	134
Synopsis.....	134
Options.....	134
See also.....	134
ibrix_version.....	135
Description.....	135
Synopsis.....	135
Options.....	135
ibrix_vg.....	136
Description.....	136
Synopsis.....	136
Options.....	137
ibrix_vs.....	138

Description.....	138
Synopsis.....	138
Options.....	138
3 Commands for Linux X9000 clients and file serving nodes.....	139
ibrix_df.....	139
Description.....	139
Synopsis.....	139
Options.....	139
Output.....	139
ibrix_lwhost	140
Description.....	140
Synopsis.....	140
Options.....	140
See also.....	141
ibrix_lwmount	142
Description.....	142
Synopsis.....	142
Options.....	142
See also.....	142
ibrix_lwumount	143
Synopsis.....	143
Options.....	143
See also.....	143
ibrix_top	144
Description.....	144
Synopsis.....	144
Options.....	145
4 Support and other resources.....	146
Contacting HP.....	146
Related information.....	146
HP websites.....	146
Subscription service.....	146
Glossary.....	147
Index.....	149

1 General CLI information

Executing X9000 Software commands

The commands in “[X9000 File Serving Software commands](#)” (page 11) must be executed on the management console host. Root privileges are required to run these commands.

The commands in “[Commands for Linux X9000 clients and file serving nodes](#)” (page 139) can be executed on file serving nodes and/or Linux X9000 clients. Any user can run these commands.

All commands are run from the working directory `/usr/local/ibrix/bin`.

Command syntax

Names. When naming physical volumes, volume groups, logical volumes, file systems, and other cluster entities, use only alphanumeric characters and the underscore (`_`). Do not use any other characters unless they are specifically required in the command syntax.

Lists. Use a comma to separate list items in a command. Spaces are not allowed. In the following command, the `-s` option takes a list of items:

```
# <installdirectory>/bin/ibrix_lv -l -s ilv1,ilv2,ilv3
```

Ranges. Many commands allow a range of values instead of an explicit list of values. To enter a range, enclose the pattern to be matched in square brackets. The range specified in the following command matches all hosts from FSHOST10 to FSHOST15.

```
# <installdirectory>/bin/ibrix_exportfs -l -h FSHOST[10-15]
```

2 X9000 File Serving Software commands

This chapter lists commands used to configure and manage X9000 File Serving Software. The commands must be run on the management console host.

ibrix_activedirectory

Configures and lists Active Directory settings.

Description

Active Directory settings must be configured on the management console to allow Windows X9000 clients to query the Active Directory server when reading or writing on a file serving node. Windows X9000 clients use the proxy user to query the Active Directory server to resolve UID/GID-to-Windows SID mappings. If mappings cannot be resolved, the user named by the `-W WIN_USER_NAME` option is displayed as the owner of the file.

Execute this command on every management console that Windows X9000 clients will need to access.

Synopsis

Configure Active Directory for static user mapping:

```
ibrix_activedirectory -S [-d DOMAIN_NAME] [-i DOMAIN_CONTROLLER_IPs]
[-u PROXY_USER] [-p PROXY_PASSWORD] [-E UID_FIELD_NAME] [-F
GID_FIELD_NAME] [-W DEFAULT_WIN_USER_NAME]
```

For the `-u` option, enter the proxy user name in the format "*domainname\username*" where *domainname* is the name of the NIS domain in Active Directory. The double quotes are required, as shown below.

```
ibrix_activedirectory -S -d fm1.hp.com -i 192.168.1.1
-u "mydomain\fusion_proxy" -p mypass -W hp_winuser
```

If Active Directory is running on Windows Server 2003 R2, the `-E` and `-F` options are required.

Configure Active Directory for automatic user mapping:

```
ibrix_activedirectory -A [-d DOMAIN_NAMES] [-L] [-W
DEFAULT_WIN_USER_NAME]
```

The `-L` option allows mapping of local users.

List Active Directory settings in effect on the management console:

```
ibrix_activedirectory -l
```

The output includes the domain name, domain server IP, proxy user name, default Windows user name, and the UID and GID field names.

Options

Option	Description
<code>-A</code>	Configures automatic user mapping.
<code>-E UID_FIELD_NAME</code>	A UID field name. On Windows Server 2003 SP2, this is <code>Read msSFU30UidNumber</code> . On Windows Server 2003 R2, it is <code>Read UidNumber</code> .
<code>-F GID_FIELD_NAME</code>	A GID field name. On Windows Server 2003 SP2 this is <code>Read msSFU30GidNumber</code> . On Windows Server 2003 R2, it is <code>Read gidNumber</code> .
<code>-L</code>	Allows local user mapping.

Option	Description
-S	Configures static user mapping.
-W <i>WIN_USER_NAME</i>	Sets the "Unknown" Windows user, who takes ownership of files when a mapping cannot be resolved.
-d <i>DOMAIN_NAME</i>	Identifies a fully qualified domain name.
-i <i>DOMAIN_CONTROLLER_IP</i>	The domain controller IP address.
-l	Displays Active Directory settings on this management console.
-p <i>PROXY_PASSWORD</i>	Sets the Active Directory proxy password.
-u <i>PROXY_USER</i>	Sets the Active Directory proxy user name, specified in the format " <i>domainname\username</i> ".
-?	Shows the usage for this command.

See also

"Windows X9000 client" in the *HP StorageWorks X9000 File Serving Software Installation Guide*

ibrix_at

Schedules X9000 Software commands.

Description

The `ibrix_at` command is used to create a schedule that is applied to another X9000 Software command. For example, `ibrix_at` can be used to create a schedule for automated snapshots.

The command has two parts, separated by a colon. The first part is a cron-like scheduling string. The following web page explains how to write the string:

<http://wiki.opensymphony.com/display/QTZ1/CronTriggers+Tutorial>

The second part of the `ibrix_at` command is the X9000 Software command to which the schedule will be applied.

Currently, `ibrix_at` can be used with `ibrix_snap` to create a snapshot schedule and then start the snapshots. The following command takes a snapshot at 11:30 a.m. on the 10th of every month. The automated snapshot strategy, specified with the `-n` option, is `monthly1`.

```
ibrix_at "0 30 11 10 * ?" : ibrix_snap -A -f ifs1 -n monthly1
```

ibrix_auth

Configures authentication for CIFS, FTP, and HTTP

Description

File serving nodes that will be hosting CIFS, FTP, or HTTP shares must be configured to authenticate with either local users and groups or your Active Directory domain.

- ❗ **IMPORTANT:** To use Local Users authentication, your cluster must use the agile management console configuration. See the *HP StorageWorks File Serving Software File System User Guide* for more information.

Synopsis

Configure Local Users authentication:

```
ibrix_auth -N [-h HOSTLIST]
```

Configure Active Directory authentication:

```
ibrix_auth -n DOMAIN_NAME -A AUTH_PROXY_USER_NAME [-P  
AUTH_PROXY_PASSWORD] [-S SETTINGLIST] [-h HOSTLIST]
```

In the command, *DOMAIN_NAME* is your Active Directory domain. *AUTH_PROXY_USER_NAME* is the name of an AD domain user (typically a Domain Administrator) having privileges to join the specified domain, and *AUTH_PROXY_PASSWORD* is the password for that account. To configure Active Directory authentication on specific nodes, specify those nodes in *HOSTLIST*.

For the *-S* option, enter the settings as *settingname=value*. Use commas to separate the settings, and enclose the list in quotation marks. If there are multiple values for a setting, enclose the values in square brackets. The users you specify must already exist. For example:

```
ibrix_auth -t -S 'share admins=[domain\user1, domain\user2,  
domain\user3]'
```

To remove a setting, enter *settingname=.*

All servers, or only the servers specified in *HOSTLIST*, will be joined to the specified Active Directory domain.

Modify authentication settings and hosts:

```
ibrix_auth -t [-S SETTINGLIST] [-h HOSTLIST]
```

List default authentication settings:

```
ibrix_auth -L
```

The settings can be specified with the *-S SETTINGLIST* option.

Display current information about authentication:

```
ibrix_auth -i [-h HOSTLIST]
```

Options

Option	Description
-A <i>AUTH_PROXY_USER_NAME</i>	Specifies the username of an account having privileges to join the Active Directory domain. This is typically a Domain Administrator.
-L	Displays default authentication settings.
-N	Configures Local Users authentication.
-P <i>AUTH_PROXY_PASSWORD</i>	Specifies the password for the proxy user account.

Option	Description
-S <i>SETTINGLIST</i>	Specifies configuration parameters for Active Directory.
-h <i>HOSTLIST</i>	A list of file serving nodes.
-n <i>DOMAIN_NAME</i>	Specifies the domain name for Active Directory authentication.
-i	Displays current information about authentication.
-t	Modifies the authentication settings.
-?	Shows the usage for this command.

ibrix_autoconnect

Manages entries in the Autoconnect table.

Description

An X9000 Software autoconnection enables NFS clients to mount file systems automatically whenever they are accessed. At the same time, Autoconnect manages how these connections are distributed among file serving nodes. Autoconnect uses the Linux `automount` daemon; working familiarity with `automount` is recommended.

Autoconnect accesses a user-edited script that directs NFS client file requests to the management console, where they are checked against the database and matched to a mount string. The mount string and any mount options are returned to the client along with the name of the file serving node that the client should use for the mount.

Mountpoints are stored in the Autoconnect table in the configuration database. Each mountpoint is described by a user-defined identifier (or key), the file system to mount, and any mount options. Autoconnections are added to or deleted from the database using the `ibrix_autoconnect -A` and `-D` commands. The `ibrix_autoconnect -l` command displays current autoconnect entries.

When setting up Autoconnect, you first need to add entries to the Autoconnect table as described here. Next, customize an Autoconnect script (either a script provided with the X9000 Software or a custom script) and then edit the `auto.master` file on NFS clients.

Synopsis

Add an entry to the Autoconnect table:

```
ibrix_autoconnect -A -k KEY -f FSNAME [-o OPTIONS]
```

KEY is a user-defined value, *FSNAME* is the file-system name, and *OPTIONS* identify NFS mount options as listed in the `linux mount` man page.

Delete entries in the Autoconnect table:

```
ibrix_autoconnect -D -k KEYLIST
```

The command deletes entries matching the keys in *KEYLIST*.

List entries in the Autoconnect table:

```
ibrix_autoconnect -l
```

Options

Option	Description
-A	Adds a mountpoint to the Autoconnect table.
-D	Deletes a mountpoint from the Autoconnect table.
-f <i>FSNAME</i>	One or more user-defined keys that represent a mountpoint.
-k <i>KEY</i> or <i>KEYLIST</i>	A key or list of keys.
-l	Lists Autoconnect table entries.
-o <i>OPTIONS</i>	Linux mount options as defined on the <code>mount</code> man page.
-?	Shows the usage for this command.

See also

`automount`, `mount`

ibrix_certificate

Manages SSL certificates.

Description

Servers accepting FTPS and HTTPS connections typically provide an SSL certificate that verifies the identity and owner of the web site being accessed. You can add your existing certificates to the cluster, enabling file serving nodes to present the appropriate certificate to FTPS and HTTPS clients. X9000 Software supports PEM certificates.

When you add an SSL certificate file to the cluster, the file must contain both the certificate contents (the .crt file) and the private key (the .key file). The certificate file must use PEM encoding, must include the headers and footers from the .crt and .key files, and cannot contain any extra spaces. See the *HP StorageWorks X9000 File Serving Software File System User Guide* for details about creating certificate files in the format required by X9000 Software.

When you configure the FTP share or the HTTP vhost, you can select the certificate to be used for FTPS or HTTPS.

Synopsis

Add a certificate:

```
ibrix_certificate -a -c CERTNAME -p CERTPATH
```

For example:

```
# ibrix_certificate -a -c mycert -p  
/usr/local/ibrix/httpd/conf/mycert.crt
```

Run the command from the active management console. To add a certificate for a different node, copy that certificate to the active management console and then add it to the cluster. For example, if node ib87 is hosting the active management console and you have generated a certificate for node ib86, copy the certificate to ib87:

```
scp server.pem ib87/tmp
```

Then, on node ib87, add the certificate to the cluster:

```
ibrix_certificate -a -c cert86 -p /tmp/server.pem
```

Delete a certificate:

```
ibrix_certificate -d -c CERTNAME
```

Display information about certificates:

```
ibrix_certificate -i [-c CERTNAME]
```

Export a certificate:

```
ibrix_certificate -e -c CERTNAME
```

This command displays the contents of the certificate. You can then copy and save the contents for future use.

Options

Option	Description
-a	Adds a certificate.
-c CERTNAME	Specifies the name of a certificate.
-d	Deletes a certificate.
-e	Exports a certificate.

Option	Description
-i	Shows information about a certificate.
-p <i>CERTPATH</i>	Specifies the location of the certificate file.

NOTE: The `-m` and `-S SETTINGLIST` options shown in the command usage message are not currently supported.

See also

`ibrix_ftpshare`, `ibrix_httpvhost`

ibrix_cfrjob

Starts and stops the replication process and monitors the state of remote replication activity.

Description

NOTE: Remote replication must be configured before replication jobs can be started.

Remote replication provides a transparent method to replicate changes in a source file system on one cluster to a target file system on either the same cluster or a second cluster. The remote replication service has two modes: continuous and run-once. For detailed information about these modes, see the *HP StorageWorks File Serving Software File System User Guide*.

Replication is started and stopped from the source cluster with `ibrix_cfrjob -s`. When a replication job is started for a file system, `ibrix_cfrjob` dynamically generates a new job with a unique job id for the replication instance.

Use the `ibrix_cfrjob -o` option for the run-once mode to synchronize single directories or entire file systems on the source and target in a single pass.

Multiple run-once jobs can be started simultaneously. The jobs terminate on their own when the synchronization is completed. Any job can be stopped manually if necessary.

While replication jobs ensure that file systems remain synchronized, run-once jobs can experience small variations if there is I/O activity on the source during the job. For example, if the source file `foo.txt` is written on the target and is then edited or deleted while the run-once job is still active, the changes to `foo.txt` are not made on the target. The `foo.txt` file on the target remains as written by the run-once job. When the job is completed, the directories containing `foo.txt` are not fully synchronized.

NOTE: If a source or target file system is unmounted and `ibrix_fsck` is run to check it, a full resynchronization is recommended.

Synopsis

Start a remote replication task to a remote cluster:

```
ibrix_cfrjob -s -f SRC_FSNAME [-o [-S SRC_DIR] ] -C TGT_CLUSTERNAME -F  
TGT_FSNAME [-P TGTDIR]
```

The `-f` option specifies the source file system to be replicated. The `-C` option specifies the target cluster. If you are replicating to a directory on the target, `-P` specifies the target directory. If the `-P` option is not used, the mount point of the target filesystem will be used as the root of the replicated data.

Use the `-o` option for run-once jobs. This option can be used to synchronize single directories or entire file systems on the source and target in a single pass. If you do not specify a source directory with the `-S` option, the replication starts at the root of the file system. The run-once job terminates after the replication is complete; however, the job can be stopped manually, if necessary.

Start an intra-cluster remote replication task:

```
ibrix_cfrjob -s -f SRC_FSNAME [-o [-S SRCDIR]] -F TGT_FSNAME [-P TGTDIR]
```

The command starts a continuous or run-once intra-cluster replication task for file system `SRC_FSNAME`. The `-F` option specifies the name of the target file system (the default is the same as the source file system). The `-P` option specifies the target directory under the target file system (the default is the root of the file system).

Use the `-o` option to start a run-once job. The `-S` option specifies a directory under the source file system to synchronize with the target directory.

Start a run-once directory replication task:

```
ibrix_cfrjob -s -f SRC_FSNAME -o -S SRCDIR -P TGTDIR
```

The `-S` option specifies the directory under the source file system to synchronize with the target directory. The `-P` option specifies the target directory.

Stop a remote replication task:

```
ibrix_cfrjob -k -n TASKID
```

Use `ibrix_task -l` to obtain the appropriate ID.

Pause a remote replication task:

```
ibrix_cfrjob -p -n TASKID
```

Use `ibrix_task -l` to obtain the appropriate ID.

Resume a remote replication task:

```
ibrix_cfrjob -r -n TASKID
```

Use `ibrix_task -l` to obtain the appropriate ID.

List all running and stopped continuous replication jobs in the cluster:

```
ibrix_cfrjob -l [-f SRC_FSNAME] [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

You can limit the output to a specific file system or host name.

Print detailed information about continuous and run-once replication tasks:

```
ibrix_cfrjob -i [-f SRC_FSNAME] [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

You can limit the output to a specific file system or host name. The display shows the status of jobs on each node, as well as job summary statistics (number of files in the queue, number of files processed). The query also indicates whether scanning is in progress on a given file serving node and lists any error conditions.

Print detailed information about continuous replication tasks that match the specified task IDs:

```
ibrix_cfrjob -i -n TASKIDS [ [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

Use `-h` to limit the output to the specified host.

View replication status and activity:

```
ibrix_cfrjob -i
```

If the command is executed from the source side, that management console answers the query with information gathered from the `ibrcfrd` daemons on the source-side file serving nodes.

Options

Option	Description
<code>-C SRC_CLUSTERNAME</code>	The remote cluster name.
<code>-F TGT_FSNAME</code>	The target file-system name.
<code>-P TGTDIR</code>	A directory under the target file system.
<code>-S SRC_DIR</code>	Identifies the source directory where replication starts. This option is valid only when used within the <code>-o</code> option, where <code>SRC_DIR</code> specifies a particular directory. The target directory pathname is assumed to be the same as that of the source path.
<code>-f SRC_FSNAME</code>	The source file-system name.
<code>-h HOSTNAME</code>	The name of the file serving node.
<code>-i</code>	Displays detailed information about continuous and run-once replication tasks.
<code>-k</code>	Stops a remote replication task.
<code>-l</code>	Lists all running and stopped continuous replication jobs.
<code>-n TASKID</code>	The ID of a continuous replication or run-once task.

Option	Description
-o	Identifies this job as a run-once action that synchronizes a particular directory.
-p	Pauses the specified remote replication task.
-r	Resumes the specified remote replication task.
-s	Starts a remote replication task for the specified file system.
-?	Shows the usage for this command.

ibrix_cifs

Manages CIFS shares.

NOTE: Be sure to use the `ibrix_cifs` command located in `<installdirectory>/bin`. The `ibrix_cifs` command located in `/usr/local/bin/init` is used internally by X9000 Software and should not be run directly.

Description

CIFS is the file-sharing protocol used in Windows-based networks. Use the `ibrix_cifs` command to configure CIFS shares. The maximum number of shares is 3000.

Access-based enumeration can be applied to a CIFS share. When this feature is in effect, users can see only the files and folders to which they have been allowed access on the CIFS share.

Synopsis

Add a CIFS share:

```
ibrix_cifs -a -f FSNAME -s SHARENAME -p SHAREPATH [-S SETTINGLIST] [-h HOSTLIST]
```

The share can be added at the specified path on all file serving nodes or only the nodes specified in `HOSTLIST`. To specify a setting, enter `settingname=value`. Use commas to separate the settings, and enclose the list in quotation marks (for example, "read only=yes,access based enumeration=true").

Modify settings on a CIFS share:

```
ibrix_cifs -m -s SHARENAME [-S SETTINGLIST] [-h HOSTLIST]
```

The share can be modified on all file serving nodes or only the nodes specified in `HOSTLIST`. To modify CIFS settings, include the `-S SETTINGLIST` option and supply the new values. To delete a setting, enter `settingname=` with no value (for example, `read only=`). To list the valid CIFS share settings, use `ibrix_cifs -L`.

Delete a CIFS share:

```
ibrix_cifs -d -s SHARENAME [-h HOSTLIST]
```

The share can be deleted from either all file serving nodes or only the nodes specified in `HOSTLIST`.

Delete all CIFS shares associated with the specified file system:

```
ibrix_cifs -d -f FSNAME
```

Display information about CIFS shares:

```
ibrix_cifs -i [-h HOSTLIST]
```

You can list information for all file serving nodes or only the nodes specified in `HOSTLIST`.

List valid CIFS share settings:

```
ibrix_cifs -L
```

Options

Option	Description
-L	Lists valid CIFS share settings.
-S SETTINGLIST	Specifies CIFS settings. Use commas to separate the settings, and enclose the list in quotation marks (for example, "read only=yes,access based enumeration=true").
-a	Adds a CIFS share.

Option	Description
-d	Deletes a CIFS share.
-f <i>FSNAME</i>	A file system.
-h <i>HOSTLIST</i>	One or more file serving nodes.
-i	Displays information about CIFS shares.
-m	Modifies a share.
-p <i>SHAREPATH</i>	The path to a shared directory.
-s <i>SHARENAME</i>	The name of a shared directory.
-?	Shows the usage for this command.

See also

`ibrix_cifsconfig`

ibrix_cifsconfig

Configures global settings for CIFS.

Description

The `ibrix_cifsconfig` command configures global settings for CIFS operations. The settings apply to all CIFS shares configured in the cluster.

Synopsis

Configure CIFS settings:

```
ibrix_cifsconfig -t [-S SETTINGLIST] [-h HOSTLIST]
```

For the `-S SETTINGLIST` option, enclose the list of settings in quotation marks, and use commas to separate the settings. For example, the following command sets SMB signing to enabled and required:

```
ibrix_cifsconfig -t -S "smb signing enabled=1,smb signing required=1"
```

To disable SMB signing, use this command:

```
ibrix_cifsconfig -t -S "smb signing enabled=,smb signing required="
```

The next example enables RFC2307, which is the protocol that enables Linux static user mapping with Active Directory:

```
ibrix_cifsconfig -t -S "rfc2307_support=rfc2307"
```

To disable RFC2307, use this command:

```
ibrix_cifsconfig -t -S "rfc2307_support=unprovisioned"
```

To see other supported settings, use the `ibrix_cifsconfig -L` command.



IMPORTANT: After making configuration changes with the `ibrix_cifsconfig -t -S` command, use the following command to restart the CIFS services on all nodes affected by the change.

```
ibrix_server -s -t cifs -c restart [-h SERVERLIST]
```

Clients will experience a temporary interruption in service during the restart.

Display current information about global CIFS settings:

```
ibrix_cifsconfig -i [-h HOSTLIST]
```

List default global CIFS configuration settings:

```
ibrix_cifsconfig -L
```

Options

Option	Description
-L	Lists default global CIFS configuration settings.
-S <i>SETTINGLIST</i>	Specifies global CIFS settings. Use commas to separate the settings, and enclose the list in quotation marks (for example, "smb_signing_enabled=1,smb_signing_required=0").
-h	Specifies one or more file serving nodes.
-i	Displays current information about global CIFS settings.
-t	Configures global CIFS settings.
-?	Shows the usage for this command.

See also

`ibrix_cifs`

ibrix_client

Performs management and control procedures for X9000 clients.

Description

`ibrix_client` enables the following X9000 client procedures:

- Registers a Linux X9000 client with a management console. Clients that are to communicate with multiple management consoles must be registered on each console. (To register a Windows X9000 client, use the Windows X9000 client GUI.)
- Sets a preferred user interface for an X9000 client. (To prefer a user interface for a hostgroup, use `ibrix_hostgroup`.)
- Unprefers a user interface for an X9000 client. (To unprefer an interface for a hostgroup, use `ibrix_hostgroup`.)
- Lists client information.
- Deletes X9000 clients from the configuration database.

X9000 clients use the cluster interface by default. It is not necessary to prefer a user interface for NFS or CIFS clients because they use client-side information to select a user interface when they mount a file system.

When a user interface has been preferred for traffic from a source host to a destination host, traffic in the reverse direction remains defaulted to the cluster interface.

The network interface preference for an X9000 client is stored in the management console. When X9000 Software services start on the client, the client queries the management console for its preferred network interface and then starts using it. If X9000 Software services are already running on a client, you can force the client to query the management console. To do this, run `ibrix_client` or `ibrix_lwhost --a`, or reboot the client.

Synopsis

Register a Linux X9000 client at an IP address:

```
ibrix_client -a -h CLIENT -e IPADDRESS
```

Execute this command on every management console with which the client will communicate.

Delete X9000 clients from the configuration database:

```
ibrix_client -d -h CLIENTLIST
```

Display operational and configuration information for X9000 clients:

```
ibrix_client -i [-h CLIENTLIST]
```

To see all clients, omit the `-h` option.

List X9000 client host names, IP addresses, and IDs:

```
ibrix_client -l [-h CLIENTLIST] [-v]
```

To see all clients, omit the `-h` option. To also list IAD, kernel, and file-system version information, include the `-v` option.

Set a preferred network interface for an X9000 client:

```
ibrix_client -n -h SRCCLIENT -A DESTSERVER/IFNAME
```

The command sets interface `IFNAME` for traffic from source client `SRCCLIENT` to destination server `DESTSERVER`. (To prefer a network interface for a hostgroup, use `ibrix_hostgroup`.)

Delete the established network interface preference for an X9000 client:

```
ibrix_client -n -h SRCHOST -D DESTSERVER
```

When the command has completed, the default cluster interface will be used to communicate from the client to the *DESTSERVER*.

Options

Option	Description
-A	Assigns a preferred NIC.
-D <i>DESTSERVER</i>	Deletes the network interface preference established between the source client and the destination server.
-a	Adds an X9000 client.
-d	Deletes a client from the configuration database.
-e <i>IPADDRESS</i>	Specifies the IP address for a client.
-h <i>CLIENTLIST</i>	One or more clients (specify as <i>CLIENT1</i> , <i>CLIENT2</i> , <i>CLIENT3</i> , . . .).
-i	Provides detailed information for one or more clients.
-l	Lists clients.
-n	Sets a preferred network interface.
-v	When used with -l, reports version information for the file system, IAD, and kernel.
-?	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_lwhost`

ibrix_cluster

Registers source and target clusters for remote replication.

Description

A cluster must be registered before other remote replication commands can use it as a target.

For each remote replication pair, cluster A must be registered with cluster B, and cluster B must be separately registered with cluster A. Run this command on both management consoles, so each cluster knows about its remote partner. You will need to run the command only once per source or target.

Synopsis

Register source and target clusters:

```
ibrix_cluster -r -C CLUSTERNAME -H REMOTE_FM_HOST
```

CLUSTERNAME is the name of the management console for a cluster. If the cluster is using an agile management console configuration, specify the `clusterName` displayed by the `ibrix_fm_tune -l` command, and enter the IP address of the cluster VIF.

Deregister a remote replication cluster:

```
ibrix_cluster -d -C CLUSTERNAME
```

List clusters registered on this management console:

```
ibrix_cluster -l
```

Options

Option	Description
-d	Deregisters a remote replication cluster.
-l	Lists clusters registered on this management console.
-r	Registers source and target clusters.
-C <i>CLUSTER_NAME</i>	Name of the target cluster being registered or deregistered.
-H <i>REMOTE_FM_HOST</i>	Name or IP address of the host where the target cluster's management console is running.
-?	Shows the usage for this command.

See also

`ibrix_cfrjob`, `ibrix_exportcfr`

ibrix_clusterconfig

Sets or displays cluster configuration parameters.

Description

Configures cluster configuration settings such as the default gateway and time zone. The settings are used during cluster setup and can also be used to update the configuration of existing systems.

Synopsis

Show cluster-wide configuration parameters:

```
ibrix_clusterconfig -i -P
```

Set cluster configuration parameters:

```
ibrix_clusterconfig -c -P PARAM1=VALUE1 [, ... , PARAMn=VALUEn]
```

Show IP address to hostname mappings:

```
ibrix_clusterconfig -i -A
```

Assign the list of host names to the specified address:

```
ibrix_clusterconfig -c -A ADDRESS -h HOSTNAME1 [, ... , HOSTNAMEn]
```

Remove the specified address from the mapping table:

```
ibrix_clusterconfig -d -A ADDRESS
```

Show the list of configured NTP servers:

```
ibrix_clusterconfig -i -N
```

Set the list of NTP servers:

```
ibrix_clusterconfig -c -N SERVER1 [, ... , SERVERn]
```

Show cluster DNS settings:

```
ibrix_clusterconfig -i -S
```

Sets the DNS servers and domains:

```
ibrix_clusterconfig -c -S SERVER1 [, ... , SERVERn] -D DOMAIN1 [, ... , DOMAINn]
```

Show network configurations:

```
ibrix_clusterconfig -i -W
```

Create or update a network configuration:

```
ibrix_clusterconfig -c -W DEV -Y CATEGORY [-M NETMASK] [-T SETUP] [-R RANGE] [-V SLAVES] [-o OPTIONS]
```

If *SETUP* is bond, slaves must be defined. *RANGE* is a comma-delimited list of ranges in the form [*nnn.nnn.nnn.nnn*-*nnn.nnn.nnn.nnn*]. *OPTIONS* is a list of *option=value* pairs, delimited by commas.

Remove a network configuration:

```
ibrix_clusterconfig -d -W DEV
```

Erase all addresses allocated for a network:

```
ibrix_clusterconfig -e -L DEV
```

Show allocated addresses for a network:

```
ibrix_clusterconfig -i -L DEV
```

Create or update an allocated address:

```
ibrix_clusterconfig -c -L DEV -I ADDRESS [-C MAC] [-h HOSTNAME]
```

Remove an allocated address:

```
ibrix_clusterconfig -d -L DEV -I ADDRESS
```

Show OS parameters and tunings:

```
ibrix_clusterconfig -i -O
```

Create or update an OS parameter:

```
ibrix_clusterconfig -c -O KEY=VALUE, [...KEY=VALUE]
```

Delete/reset an OS parameter:

```
ibrix_clusterconfig -d -O KEY
```

List OS parameters and defaults:

```
ibrix_clusterconfig -Q
```

Synchronize parameters on file serving nodes:

```
ibrix_clusterconfig -s [-P] [-O] [-o keyword] [-h HOSTLIST]
```

Use `-P` to synchronize cluster parameters such as DNS and NTP. Use `-O` to synchronize OS tunings.

Options

Option	Description
-A	Host address aliases.
-C	MAC address.
-D DOMAIN1 [, ... , DOMAINn]	Sets DNS domains.
-I ADDRESS	IP address [nnn . nnn . nnn . nnn].
-L DEV	Allocated addresses.
-M NETMASK	Netmask [nnn . nnn . nnn . nnn].
-N SERVER1 [, ... , SERVERn]	NTP servers.
-O KEY=VALUE, [...KEY=VALUE]	OS parameters and tunings.
-P	Named configuration parameters.
-Q	Queries OS parameters and tunings.
-R RANGE	Address range [nnn . nnn . nnn . nnn - nnn . nnn . nnn . nnn].
-S SERVER1 [, ... , SERVERn]	DNS servers.
-T SETUP	Network setup type.
-V SLAVES	Bonded slave devices (DEV1, DEV2, ...).
-W DEV	Preconfigured networks.
-Y CATEGORY	Network category (cluster or user).
-c	Changes the stored configuration.
-d	Deletes the stored configuration.
-e	Erases address allocations.
-h	A list of one or more hosts (HOST1, HOST2, HOST3, ...).
-i	Shows the stored configuration.
-o keyword	Option list.
-s	Synchronizes settings on file serving nodes.
?	Shows the usage for this command.

ibrix_dbck

Ensures that information maps on hosts are consistent with the configuration database.

NOTE: `ibrix_dbck` should be used only under the direction of HP Support.

Description

Hosts (file serving nodes and X9000 clients) maintain information about a file system based on information that they obtain from the configuration database. If a host's information about a file system becomes outdated, the host cannot access the file system or communicate with other hosts about the file system.

Run `ibrix_health` regularly to keep this information up to date. If the information becomes outdated on a host, execute `ibrix_dbck -o` to resynchronize the information on the hosts with the configuration database.

Synopsis

Update file-system information on hosts:

```
ibrix_dbck -o -f FSNAME [-h HOSTLIST]
```

To update all hosts, omit the `-h HOSTLIST` option.

Options

Option	Description
<code>-f <i>FSNAME</i></code>	A file system.
<code>-h <i>HOSTLIST</i></code>	A list of one or more file serving nodes or X9000 clients.
<code>-o</code>	Synchronizes cluster information.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_health`

ibrix_edquota

Sets and manages file system quotas for users, groups, and directory trees.

Description

The `ibrix_edquota` command sets, lists, and deletes storage quotas for users or groups on a specific file system. It also sets directory tree quotas, which limit the amount of storage and the number of files that can be created on a file system located at a specific directory tree. By default, quotas are not set.

The recommended order for handling quotas is as follows:

1. Enable quotas when creating the file system. Quotas must be enabled before quota limits can be set. (You can also enable quotas on an existing file system. First, unmount the file system, and then enable quotas with `ibrix_fs -q -E -f FSNAME`.)
2. Mount the file system.
3. For users and groups, set quota limits with `ibrix_edquota -s`.
4. For directory trees, create a directory tree identifier with the `ibrix_fs_ops` command, and then assign quota limits to the directory tree with `ibrix_edquota -s`.

A quota is delimited by hard and soft storage limits defined either in megabytes of storage or as a number of files. The hard limit is the maximum storage (in terms of file size and number of files) allotted to a user or group. The soft limit specifies the number of megabytes or files that, when reached, starts a countdown timer that runs until the hard storage limit is reached or seven days elapse, whichever happens first. When the timer stops for either reason, the user or group cannot store any more data and the system issues `quota exceeded` messages at each write attempt.

When setting quotas, no UID or GID can exceed 2,147,483,647. Setting user quotas to zero removes the quotas.

Quotas information can be imported into the system from a file. The existing quotas information can also be exported to a file. See the *HP StorageWorks X9000 File Serving Software File System User Guide* for information about the quotas file.

NOTE: When a new NIS user is added, you will need to restart the Fusionmanager services before assigning quotas to the user:

```
/etc/init.d/ibrix_fusionmanager restart
```

Synopsis

Set a quota for a single user on the specified file system:

```
ibrix_edquota -s -u "USER" -f FSNAME [-M SOFT_MEGABYTES] [-m  
HARD_MEGABYTES] [-I SOFT_FILES] [-i HARD_FILES]
```

Enclose the user name in single or double quotation marks.

Set a quota for a single group on the specified file system:

```
ibrix_edquota -s -g "GROUP" -f FSNAME [-M SOFT_MEGABYTES] [-m  
HARD_MEGABYTES] [-I SOFT_FILES] [-i HARD_FILES]
```

Enclose the group name in single or double quotation marks.

Assign usage limits to a directory tree quota:

```
ibrix_edquota -s -d NAME -f FSNAME [-M SOFT_MEGABYTES] [-m  
HARD_MEGABYTES] [-I SOFT_FILES] [-i HARD_FILES]
```

Import quotas information from a file:

```
ibrix_edquota -t -p PATH [-f FSNAME]
```


Export the existing quotas information to a file:

```
ibrix_edquota -e -p PATH [-f FSNAME]
```

Delete a user quota:

```
ibrix_edquota -D -u UID [-f FSNAME]
```

To delete the quota on a specific file system, include the `-f` option. Omit this option to delete the quota on all file systems.

Delete a group quota:

```
ibrix_edquota -D -g GID [-f FSNAME]
```

To delete the quota on a specific file system, include the `-f` option. Omit this option to delete the quota on all file systems.

Delete the usage limits for a directory tree quota:

```
ibrix_edquota -D -d NAME -f FSNAME
```

To delete the quota on a specific file system, include the `-f` option. Omit this option to delete the quota on all file systems.

List quota limits and usage for a specific user:

```
ibrix_edquota -l -u UID [-f FSNAME]
```

List quota limits and usage for all users:

```
ibrix_edquota -l -U [-f FSNAME]
```

List quota limits and usage for a specific group:

```
ibrix_edquota -l -g GID [-f FSNAME]
```

List quota limits and usage for all groups:

```
ibrix_edquota -l -G [-f FSNAME]
```

List usage limits for a specific directory tree quota:

```
ibrix_edquota -l -d NAME [-f FSNAME]
```

List usage limits for all directory tree quotas or for quotas on a specific file system:

```
ibrix_edquota -l -T [-f FSNAME]
```

Options

Option	Description
<code>-D</code>	Deletes user, group, or directory tree limits.
<code>-G</code>	Specifies all groups with quota settings.
<code>-I SOFT_FILES</code>	Sets the soft limit on the number of files.
<code>-M SOFT_MEGABYTES</code>	Sets the soft limit on the megabyte allocation.
<code>-T</code>	All directory trees.
<code>-U</code>	Specifies all users with quota settings.
<code>-d NAME</code>	Specifies the name of a directory tree quota.
<code>-e PATH</code>	Exports the existing quotas information to a file.
<code>-f FSNAME</code>	A file system.
<code>-g "GROUP" or GID</code>	A group or group ID. Enclose the group name in single or double quotation marks. The GID value cannot exceed 2,147,483,647.
<code>-i HARD_FILES</code>	Sets the hard limit on the number of files. The default value is 0.

Option	Description
-l	Lists quota limits.
-m <i>HARD_MEGABYTES</i>	Sets the hard limit for megabyte allocation. The default value is 0.
-s	Sets quota limits.
-t <i>PATH</i>	Imports quotas information from a file.
-u " <i>USER</i> " or <i>UID</i>	A user name or user ID. Enclose the user name in single or double quotation marks. The UID value cannot exceed 2,147,483,647.
-?	Shows the usage for this command.

See also

`ibrix_fs`, `ibrix_fs_ops`

ibrix_event

Sets up email and SNMP notifications of cluster events, displays event messages, and removes event messages from the events table.

Description

Use `ibrix_event` to set up two types of event notifications:

- Email notification, which sends to one or more addresses an email containing the event, the originating server, the severity level, a timestamp, an autogenerated number, and the management console name.
- SNMP notification, which sends an SNMP trap to one or more trapsinks.

X9000 Software events fall into categories based on level of severity:

- **Alerts:** Disruptive events that can result in loss of access to file-system data. For example, a segment is unavailable or a server is unreachable.
- **Warnings:** Potentially disruptive conditions in which file-system access is not lost, but if the situation is not addressed it can escalate to an alert condition. For example, reaching a very high server CPU utilization or nearing a quota limit.
- **Information:** Events occurring under normal or non-threatening conditions that change the cluster. For example, creating a segment or mounting a file system.

To set up email notifications, first associate events and email recipients with `ibrix_event -c` and then configure settings and initiate the notification process with `ibrix_event -m`. Sender and recipient email addresses must be known to the SMTP server.

To set up SNMP notifications, first associate events and trapsinks with `ibrix_event -c` and then configure notification settings and initiate the notification process with `ibrix_snmptrap -c`.

Synopsis

Configure email notification settings:

```
ibrix_event -m on|off -s SMTP -f from [-r reply-to] [-t Subject]
```

When `-m on` is specified, all options are required.

Associate events with email addresses:

```
ibrix_event -c [-e {ALERT|WARN|INFO|EVENTLIST}] -m EMAILLIST
```

The `-e` option associates events with event types. You can enter any combination of types or a list of messages. Omit the `-e` option to associate all event types. For example, to associate all Alert events and two Info events to `admin@mycompany.com`:

```
ibrix_event -c -e ALERT,server.registered,filesystem.created  
-m admin@mycompany.com
```

Associate events with trapsinks:

```
ibrix_event -c -y SNMP [-e {ALERT|WARN|INFO|EVENTLIST}] -m TRAPSINKLIST
```

Specify the hostname or IP address of the trapsink. The `-e` option associates events with event types. Omit the `-e` option to associate all event types. For example, to associate all Alert events and two Info events with a trapsink at IP address `192.168.2.32`, enter:

```
ibrix_event -c -y SNMP -e ALERT,server.registered,  
filesystem.created -m 192.168.2.32
```

Test the supplied email address:

```
ibrix_event -u -n EMAILADDRESS
```

Before sending a test message to the email address, notifications must be turned on. If the address is valid, the command signals success and sends an email containing the settings to the recipient. If the address is not valid, the command returns an `address failed` exception.

Remove events from the events table:

```
ibrix_event -p [-o DAYS_COUNT]
```

Events are removed starting with the oldest events. The default is to remove the oldest seven days of messages. To change the number of days, include the `-o DAYS_COUNT` option.

List all events, or only events of the specified type:

```
ibrix_event -q [-e ALERT|WARN|INFO]
```

Lists email settings and all configured notifications:

```
ibrix_event -L [-y EMAIL|SNMP]
```

Include `-y` with either `EMAIL` or `SNMP` to list settings for only those types of events.

Lists posted events in reverse chronological order:

```
ibrix_event -l [-f FILESYSTEM] [-h SERVER] [-v {ALERT|WARN|INFO}] [-s STARTTIME] [-e ENDTIME] [-n EVENTS_COUNT] [-a]
```

To list events in ascending order, include the `-a` option. The command displays the most recent 100 messages by default. This can be increased or decreased with the `-n COUNT` option. The other options act as filters to return records associated with the specified file system, server, alert level, and start or end time.

Options

Option	Description
<code>-L</code>	Lists event notifications and settings.
<code>-a</code>	Sorts in ascending order.
<code>-c</code>	Creates email notifications.
<code>-d</code>	Deletes email notifications.
<code>-e ALERT WARN INFO EVENTLIST</code>	The ALERT, WARN or INFO cluster event notifications or a list of specific notification types. ALERT, WARN and INFO are keywords.
<code>-f FROM or FILENAME</code>	An email "From" address or, if used with <code>-v</code> , the name of the file where a real-time event stream is saved in CSV format.
<code>-h SERVER</code>	A file serving node.
<code>-l</code>	Lists event notifications and settings.
<code>-m on off or EMAILLIST TRAPSINKLIST</code>	Turns event notification on or off, or if used with <code>-c</code> or <code>-d</code> , specifies a list of one or more emails or trapsinks to be created or deleted.
<code>-n EMAILADDRESS or EVENTS_COUNT</code>	When used with <code>-u</code> , specifies a destination address for a test email message. When used with <code>-l</code> , specifies the number of event messages to include, in reverse order, in the output.
<code>-o DAYS_COUNT</code>	When removing messages from the events table with <code>-p</code> , sets the number of days worth of messages to remove.
<code>-p</code>	Removes events from the events table, starting with the oldest message. The default is to remove seven days of messages.
<code>-q</code>	Lists all configurable events.
<code>-r REPLY-TO</code>	Sets the email Reply-to address.
<code>-s SMTPADDR</code>	Sets the IP address of the SMTP server.

Option	Description
-t <i>SUBJECT</i>	Specifies an email subject string that is prefixed to the automatically generated email Subject line.
-u	Sends a test email.
-v {ALERT WARN INFO}	An event level.
-y EMAIL SNMP	Specifies an event notification type. The valid values are EMAIL (the default) and SNMP.
-?	Shows the usage for this command.

See also

`ibrix_snmptrap`

“Setting up email notification of cluster events” and “Setting up SNMP notifications” in the administration guide for your system.

ibrix_exportcfr

Configures a target (a file system or directory) for remote replication.

Description

This command exports the target file system. Before replication can take place, you must create a mapping between the source cluster and the target file system that receives the replicated data. This mapping ensures that only the specified source cluster can write to the target file system.

NOTE: You do not need to export the file system when configuring intracluster replication.

Synopsis

Export a target file system:

```
ibrix_exportcfr -f FSNAME [-p directory] -C REMOTE_CLUSTER [-P]
```

The *FSNAME* is the target file system to be exported. The *-p* option can be used to export a directory located under the root of the specified file system. The *-C REMOTE_CLUSTER* option specifies the source cluster (clusterName) containing the file system to be replicated. Include the *-P* option if you do not want to set the default host and NIC preferences. The following example exports file system *ifs1* on the target cluster. Source cluster *dev3.company.com* contains the file system that will be replicated to *ifs1* on the target:

```
# ibrix_exportcfr -f ifs1 -C dev3.company.com
```

Unexport a file system (or a directory in that file system) for remote replication:

```
ibrix_exportcfr -U -f FSNAME [-p DIRECTORY]
```

Remote replication will fail if the target file system is unexported. If replication jobs are using the export, wait for them to finish or cancel them before unexporting the file system.

List remote replication exports:

```
ibrix_exportcfr -l
```

Options

Option	Description
<i>-C REMOTE_CLUSTER</i>	The name of the source cluster containing the file system to be replicated.
<i>-P</i>	Skips setting the default preferences.
<i>-U</i>	Unexports a file system for remote replication.
<i>-f TARGET_FSNAME</i>	The file system that is the target for the replication.
<i>-p DIRECTORY</i>	Exports a directory within the target file system.
<i>-l</i>	Lists all exported file systems.
<i>-?</i>	Shows the usage for this command.

See also

`ibrix_cluster`

ibrix_exportcfrpreference

Identifies the preferred hosts for replication storage and a preferred NIC for replication traffic.

Description

If you are replicating a file system to a remote cluster, you can select the hosts that will store replication data and the network NIC that will carry replication traffic on the cluster. The default preferences are:

- Use all hosts that have the file system mounted.
- Use the cluster NIC on each host.

If you prefer to select the hosts and, optionally, the network NIC that will be used on the remote cluster, use the `ibrix_exportcfrpreference` command to set your own preferences.

Synopsis

Identify the preferred hosts and, optionally, a preferred NIC:

```
ibrix_exportcfrpreference -a -f FSNAME [-p directory] -h HOSTLIST [-n NETWORK INTERFACE NAME]
```

The `-p` option specifies a directory under the root of file system *FSNAME*.

When specifying resources, note the following:

- Specify hosts by their host name or IP address. A host is any host on the target cluster that has the target file system mounted.
- Specify the network using the X9000 Software network name. Enter a valid user NIC or the cluster NIC. Setting the network preference is optional. If it is not specified, the host name (or IP) is used to determine the network.

The listed hosts receive remote replication data via the specified NIC. To increase capacity, you can expand the number of preferred hosts by executing this command again with another list of hosts.

Restore the default preferences for remote replication:

```
ibrix_exportcfrpreference -D -f FSNAME [-p directory]
```

The `-p` option specifies a directory under the root of file system *FSNAME*.

Remove a remote replication preference:

```
ibrix_exportcfrpreference -r -P PREFERENCE_ID_LIST
```

To obtain the ID for a particular preferred host, list the remote replication preferences using the `-l` option.

View preferences for remote replication:

```
ibrix_exportcfrpreference -l
```

The output lists the exported file systems and associated host preferences on this cluster. All preferred hosts/NICs are listed with a corresponding ID number that can be used in commands to remove preferences.

Options

Option	Description
<code>-D</code>	Restores the default preferences (automatically use all hosts and the cluster network interface on each host).
<code>-P <i>PREFERENCE_ID_LIST</i></code>	Preference ID.

Option	Description
-a	Identifies a preference.
-f <i>FSNAME</i>	The exported file system.
-h <i>HOSTLIST</i>	Host preferences (list of hosts used by the file system).
-l	Lists the export file systems and preferences on this cluster.
-n <i>NETWORK INTERFACE NAME</i>	Network interface preference (use the X9000 Software network interface name).
-p <i>directory</i>	Exported directory.
-r	Removes a preference.
-?	Usage for this command.

ibrix_exportfs

Maintains the list of NFS-exported file systems and specifies file-system access for X9000 clients.

Description

`ibrix_exportfs` makes local directories available for NFS clients to mount. The management console manages the table of exported file systems and distributes the information to local `/etc/exports` files. All entries are automatically re-exported to NFS clients and, optionally, to the file serving node standbys. A file system must be mounted before it can be exported. It must be unexported before can be unmounted.

The `ibrix_exportfs` command can also be used to control file-system access for X9000 clients. By default, all X9000 clients can mount a file system after a mount point has been created. To allow access for a file system (or a subdirectory of that file system) to a specific X9000 clients, use the `-c` option to create an access entry containing that information. Then, when an X9000 client attempts to mount the file system, X9000 Software will allow access only to the clients in the access entry.

Synopsis

Export a file system via NFS:

```
ibrix_exportfs -f FSNAME -h HOSTNAME -p  
CLIENT1:PATHNAME1,CLIENT2:PATHNAME2,... [-o "OPTIONS"] [-b]
```

File system *FSNAME* at mountpoint *PATHNAME* on file serving node *HOSTNAME* is exported to the specified *CLIENTS*. The accessing *CLIENT* can be a single file serving node, file serving nodes represented by a wildcard, or all hosts (`:/PATHNAME`), which is also known as world access. World access omits the client specification but not the colon, for example, `:/usr/src`. By default, the file system is exported to the NFS client's standby. To exclude the standby, include the `-b` option. The command uses the default Linux `exportfs` options unless specific options are provided.

Unexport the specified NFS exports:

```
ibrix_exportfs -U -h HOSTLIST -p CLIENT1:PATHNAME1,CLIENT2:PATHNAME2,...  
[-b]
```

The `-b` option excludes the standby.

List NFS-exported directories:

```
ibrix_exportfs -l [-h HOSTLIST]
```

To specify a range of servers in the host list, enter the pattern to match, enclosed in square brackets. Separate items in a list with commas.

Create an access entry for a file system:

```
ibrix_exportfs -c -f FSNAME -p CLIENT1:PATHNAME1,CLIENT2:PATHNAME2,...
```

By default, all X9000 clients can mount a file system after a mount point has been created. This command limits the clients that can mount the file system or a subdirectory of the file system. The `-p` option specifies the clients and the file systems (or subdirectories) that they are allowed access. *CLIENT* is the name of an X9000 client or hostgroup and */PATHNAME* is the file system or subdirectory the client will be allowed to mount.

For example, the following command gives client2 access to the subdirectory `/ifs1/tests` in the `ifs1` file system:

```
ibrix_exportfs -c -f ifs1 -p client2:/ifs1/tests
```

Remove the access entry for a client:

```
ibrix_exportfs -U -f FSNAME -p CLIENT1:PATHNAME1,CLIENT2:PATHNAME2,...
```

Options

Option	Description
-U	Unexports the named file system.
-b	Excludes the file serving node standby from the export.
-c	Specifies file-system access for an X9000 client.
-f <i>FSNAME</i>	A file system.
-h <i>HOSTNAME</i> or <i>HOSTLIST</i>	A file serving node or a list of file serving nodes exporting a file system.
-l	Lists all NFS export entries, or lists entries on a named host.
-o " <i>OPTIONS</i> "	Linux <code>exportfs</code> options. Standard NFS export options are supported. Options must be enclosed in double quotation marks (for example, -o "ro"). Do not enter an <code>FSID=</code> or <code>sync</code> option. They are automatically provided.
-p <i>CLIENT:PATHNAME</i>	The path to the exported file system and the NFS client that can access it. The NFS client can be specified by the name or IP of a file serving node, a file serving node name with a wildcard (for example, *.hp.com), or a netmask.
-?	Shows the usage for this command.

See also

Linux `exportfs` command

"Exporting a file system" in the *HP StorageWorks X9000 File Serving Software File System User Guide*

ibrix_fm

Controls an agile management console configuration.

Description

Typically, one active management console and one or more passive management consoles are installed when the cluster is installed. This is called an *agile* management console configuration.

An agile management console can be in one of the following modes:

- **active.** In this mode, the management console controls console operations. All cluster administration and configuration commands must be run from the active management console.
- **passive.** In this mode, the management console monitors the health of the active management console. If the active management console fails, X9000 Software selects one of the passive management consoles to become the active console.
- **maintenance.** In this mode, the management console does not participate in console operations. This mode is intended for operations such as manual failover of the management console, X9000 Software upgrades, and blade replacements.

Synopsis

Set the mode for a management console:

```
ibrix_fm -m [passive | maintenance]
```

Moves a management console to passive or maintenance mode. If the management console was previously the active console, X9000 Software will select a new active console. A management console currently in active mode can be moved to either passive or maintenance mode. A management console in maintenance mode can be moved only to passive mode.

Run an immediate backup of the management console configuration:

```
ibrix_fm -B
```

By default, the command saves the backup archive at `<installdirectory>/tmp/fmbbackup.zip`. The existing `fmbbackup.zip` file is overwritten.

Add virtual interface information to the management console configuration:

```
ibrix_fm -c <VIF IP address> -d <VIF device> -n <VIF netmask> -v <VIF type> [-I <local IP address>]
```

The `<VIF type>` is either `cluster` or `user`.

Register a management console and its IP address/hostname with the active management console:

```
ibrix_fm -R <management_console_name> -I <IP address> -a <cluster VIF address>
```

Unregister a management console that has been uninstalled:

```
ibrix_fm -u <management_console_name>
```

Display mode information about the management consoles in the cluster:

```
ibrix_fm -i
```

NOTE: If the management console was not installed in an agile configuration, the output will report `FusionServer: fusion manager name not set! (active, quorum is not configured)`.

List registered management consoles:

```
ibrix_fm -f
```

Options

Option	Description
-B	Backs up the management console configuration.
-I <IP address>	Specifies an IP address.
-R	Registers a management console.
-a <cluster VIF address>	Specifies the IP address of the cluster virtual interface.
-c <VIF IP address>	Adds information about a virtual interface to the management console configuration.
-d <VIF Device>	Specifies the name of the device used for the virtual interface.
-f	Lists registered management consoles.
-i	Displays mode information for management consoles.
-m [passive maintenance]	Moves a management console to passive or maintenance mode.
-n <VIF netmask>	Specifies the netmask for the virtual interface.
-u management console name	Unregisters a management console that has been uninstalled.
-v <VIF type>	Specifies the type of interface (cluster or user).
-?	Shows the usage for this command.

ibrix_fm_tune

Sets or resets configuration tuning parameters on the management console.

Description

⚠ CAUTION: This command sets or resets the configuration tuning parameters on the management console.

The default values for the management console tuning parameters are suitable for most cluster environments. Changing parameter values can alter file-system performance. HP recommends that you exercise caution before implementing any changes, or do so only under the guidance of HP technical support.

Synopsis

Set configuration options on the management console:

```
ibrix_fm_tune -S -o OPTIONLIST
```

Options are listed in [Table 1 \(page 45\)](#). *OPTIONLIST* is specified as a comma-delimited list of option/value pairs. Enter the option name as it appears in the first column in the table. For example:

```
ibrix_fm_tune -S -o healthCheckInterval=45,usageStatisticsSaveInterval=1800
```

Reset management console configuration options to their default values:

```
ibrix_fm_tune -U -n OPTIONS
```

Enter the option name as it appears in the first column in the table, and separate option names with commas, for example:

```
ibrix_fm_tune -U -n healthCheckInterval,usageStatisticsSaveInterval
```

List all management console configuration options:

```
ibrix_fm_tune -L
```

The output includes the default values and valid value ranges. The options are listed in the configuration options table.

List management console configuration options set to non-default values:

```
ibrix_fm_tune -l [-n OPTIONS]
```

The *-n* option displays only the listed options and their values.

Turn up the logging levels:

```
ibrix_fm_tune -v logginglevel -c logging component
```

Table 1 *ibrix_fm_tune* configuration options

Option name	Description
bufferOverflowMonitorInterval	The interval for the buffer overflow monitor, specified in seconds. Value is an integer. Default: 60 sec.
clusterName	The name of the cluster.
cpuThresholdSensitivityInterval	Interval in seconds between samples when the CPU utilization exceeds the <i>cpuUsageEventThreshold</i> . File serving nodes only. Range: 0-65535. Default: 300 sec.
cpuUsageEventThreshold	File serving node CPU utilization threshold. When exceeded, triggers an event. Value is a percentage. Default: 90 percent.
enableMigrationHealthCheck	Sets whether a host health check is performed during segment migration. Value is true or false. Default: false.

Table 1 ibrix_fm_tune configuration options *(continued)*

Option name	Description
enableSegmentUnavailableFailover	Sets whether the management console will implement a policy other than the default (if any segment is unavailable, then fail over) when a segment is unavailable. Value is true or false. Default: false.
filesystemSpaceUsedEventThreshold	Total file-system disk utilization threshold. When exceeded, triggers an event. Value is a percentage. Default: 90 percent.
fusionManagerPrimaryAddress	IP address.
fusionManagerSecondaryAddress	IP address.
healthCheckInterval	Interval in seconds between health checks on managed hosts. Range: 1-86400. Default: 60 sec.
maxMdDevices	The maximum number of multiple disk devices allowed. Range: 0-128. Default: 32.
maxNbdDevices	The maximum number of network block devices allowed. Range: 0-128. Default: 32 devices.
minFailbackInterval	The minimum number of seconds that must elapse before attempting another failback. Value is an integer. Default: 60 sec.
nbdPortsLimit	The maximum number of network block device ports allowed. Range: 0-128. Default: 32 ports.
nbdTimeout	The maximum network block device timeout. Range: 10-300. Default: 60 sec.
quotaMonitorLoopInterval	Quota monitor loop interval. Value is an integer. Default: 300.
segmentSpaceUsedEventThreshold	Segment storage utilization threshold. When exceeded, generates an event. Value is a percentage. Default: 90 percent.
segmentUnavailableFailoverPolicy	Management console policy to follow when a segment is unavailable. Values are 0 (if any segment is unavailable, then fail over) or 1 (if all segments are unavailable, then fail over). Default: 0.
segmentUnavailableScanInterval	Interval in seconds between management console checks for unavailable segments. Range: 0-100. Default: 30 sec.
serverPort	Port used for cluster network communication. Range: 0-65535. Default: port 1234.
skippedCyclesBeforeClientFailure	Sets the number of health reporting cycles that must fail before a client is considered to have failed. Value is an integer. Default: 120.
skippedCyclesBeforeFailure	Sets the number of health reporting cycles that must fail before a file serving node is considered to have failed. Value is an integer. Default: 5.
snapUsageMonitorInterval	Interval in seconds between Snap usage checks. Value is an integer. Default: 300 sec.
usageStatisticsSaveInterval	Interval in seconds between writes of usage statistics. Value is an integer. Default: 3600 sec.

Options

Option	Description
-L	Lists default configuration tuning options.
-S	Sets configuration tuning options.

Option	Description
-U	Resets configuration tuning options.
-c	Selects the logging component.
-l	Lists configuration tuning options.
-n <i>OPTIONS</i>	Specifies a list of comma-separated options in the format <i>option1,option2</i> and so on.
-o <i>OPTIONLIST</i>	Specifies a list of comma-separated option/value pairs in the format <i>option1=value1,option2=value2</i> and so on.
-v	Sets the logging level.
-?	Shows the usage for this command.

See also

`ibrix_fs_tune`, `ibrix_host_tune`

ibrix_fs

Creates, extends, and manages file systems.

Description

`ibrix_fs` is used to:

- Create a file system from physical or logical volumes.
- List information about file systems (including snapshot file systems).
- Extend a file system by adding new segments.
- Delete a file system.
- Migrate segments to another file serving node.
- Enable or disable export control on a file system.
- Disable 32-bit mode on a file system.
- Assign or unassign a standby server.

Synopsis

Create a file system

Before creating a file system, determine the mode for the file system and whether data tiering should be used:

- **32-bit or 64-bit mode.** If all file-system clients (NFS, CIFS, and X9000 clients) run 64-bit applications exclusively, create the file system to use 64-bit mode. If 32-bit applications are in use, create the file system to use 32-bit mode, which is the default. (A file system created in 32-bit mode can be converted to use 64-bit mode by disabling 32-bit mode.)
For 64-bit mode, there is no restriction on the maximum number of segments allowed per file system. For 32-bit mode, the maximum number of segments allowed per file system is 255. You can set the maximum number of segments allowed when you create the file system.
- **Data tiering.** A data tier is a group of one or more segments. Tiering is the ability to automatically migrate files from one tier to another within the same file system. Tiering is optional and it is turned off by default; however, you can enable it when creating or extending a file system.

To create a file system, execute one of the following `ibrix_fs -c -f` commands.

Create a file system with the specified segments:

```
ibrix_fs -c -f FSNAME -s LVLIST [-t TIERNAME] [-a] [-q] [-o  
OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME] [
```

The command creates the file system using the segments specified in *LVLIST*. The round-robin method is used to assign segments to file serving nodes.

Create a file system and assign the specified segments to the specified file serving nodes:

```
ibrix_fs -c -f FSNAME -S LV1:HOSTNAME1,LV2:HOSTNAME2,... [-a] [-q] [-o  
OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME]
```

The command creates file system *FSNAME* and assigns specific segments (*LV1*, *LV2*, ...) to specific file serving nodes (*HOSTNAME1*, *HOSTNAME2*, ...). The first logical volume listed becomes the root segment for the file system. The round-robin method is used to assign segments to file serving nodes.

Create a file system from the specified physical volumes in one step:

```
ibrix_fs -c -f FSNAME -p PVLIST [-a] [-q] [-o  
OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME]
```

The file system is created from the physical volumes in *PVLIST*. One volume group is created per physical volume. Segments are then created from the volume groups and assigned to file serving nodes in a round-robin manner.

Other options for creating a file system

The following options can be used with any of the create commands:

-a

Create a mountpoint with the same name as the file system and then mount the file system.

-q

Enable quotas on the file system.

-t *TIERNAME*

Assign the file-system segments to the specified tier.

-o compat=no

Enable 64-bit mode on the file system.

-o compat=yes,maxsecs=*n*

Create a 32-bit file system and specify the maximum number of segments to allow in the file system. The valid values for maxsecs are 15, 31, 63, 127, and 255.

If none of these options are specified, the new file system is created with quotas disabled, with no mountpoint, and with 32-bit mode compatibility enabled.

View information about file systems

Use the following commands to display details about the file systems configured on the cluster.

Display file system information:

```
ibrix_fs -i [-f FSLIST] [-x]
```

The output includes either all file systems or only the file systems in *FSLIST*, as well as their component segments. Include -x to omit segment details.

List information about all file systems on all hosts:

```
ibrix_fs -l
```

For each file system, the command notes its name, number of segments, whether it is mounted, the file-system generation, whether quotas and export control are enabled, and whether it is backward-compatible with 32-bit systems.

Extend a file system

Use any of the following commands to extend a file system.

Extend a file system using the specified logical volumes:

```
ibrix_fs -e -f FSNAME -s LVLIST [-t TIERNAME]
```

The file system is extended using the logical volumes in *LVLIST*. For example, to add logical volumes *ilv1*, *ilv2*, and *ilv3* to file system *ifs1*:

```
ibrix_fs -e -f ifs1 -s ilv[1-3]
```

To extend the file system with *LVLIST* representing tiered segments, include the -t option. If tiering rules are already defined for this file system, -t is required.

Extend a file system using the specified segment/owner pairs:

```
ibrix_fs -e -f FSNAME -S LV1:HOSTNAME1, LV2:HOSTNAME2,... [-t TIERNAME]
```

For example, to add logical volume `ilv1` (owned by `s1.mycompany.com`) to file system `ifs1`:

```
ibrix_fs -e -f ifs1 -S ilv1:s1.mycompany.com
```

To extend the file system with tiered segments, include the `-t` option. If tiering rules are already defined for this file system, `-t` is required.

Extend a file system using the specified physical volumes:

```
ibrix_fs -e -f FSNAME -p PVLIST [-t TIERNAME]
```

The file system is extended using the physical volumes in `PVLIST`. To extend the file system with tiered segments, include the `-t` option. If tiering rules are already defined for this file system, `-t` is required.

Delete a file system

Delete the specified file systems:

```
ibrix_fs -d -f FSLIST
```

Migrate segment ownership

Use the following commands to migrate ownership of segments to another file serving node. Migration does not move segments from one physical location to another. Instead, the ownership is recorded on the physical segment itself.

Migrate ownership of segments to a new host and update the source host:

```
ibrix_fs -m -f FSNAME -s LVLIST -h HOSTNAME [-M] [-F] [-N]
```

The segments in `LVLIST` on file system `FSNAME` are migrated to the new host. To force the migration, include `-M`. To skip the source host update during the migration, include `-F`. To skip host health checks, include `-N`.

Migrate ownership of segments from one host to another and update the source host:

```
ibrix_fs -m -f FSNAME -h HOSTNAME1,HOSTNAME2 [-M] [-F] [-N]
```

The segments owned by `HOSTNAME1` are migrated to `HOSTNAME2`. To force the migration, include `-M`. To skip the source host update during the migration, include `-F`. To skip host health checks, include `-N`.

Enable or disable quotas

Enable quotas on a file system:

```
ibrix_fs -q -E -f FSNAME
```

The file system must be unmounted when this command is run. The command may take some time to finish, especially if the file system is heavily populated.

Disable quotas on a file system:

```
ibrix_fs -q -D -f FSNAME
```

Quotas can be disabled on either mounted or unmounted file systems.

Mark a bad segment

Mark segment number(s) bad in the specified file system:

```
ibrix_fs -B -f FSNAME {-n BADSEGNUMLIST | -s BADLVLIST}
```

The `-n` option uses segment numbers to indicate bad segments. The `-s` option uses `lvm` names to indicate bad segments.

Enable or disable Export Control

When Export Control is enabled, by default X9000 clients have no access to the file system. The system administrator grants access to the clients by executing the `ibrix_mount` command.

Enabling Export Control does not affect file-system access by file serving nodes or any NFS/CIFS clients attached to the nodes. File serving nodes always have RW access.

Enable Export Control on a file system:

```
ibrix_fs -C -E -f FSNAME
```

Disable Export Control on a file system:

```
ibrix_fs -C -D -f FSNAME
```

Disable 32-bit mode

Disabling 32-bit mode causes the file system to be in 64-bit mode.



IMPORTANT:

Disabling 32-bit mode is a one-time operation and cannot be reversed; if clients will need to run 32-bit applications in the future, do not enable 64-bit mode.

Disable 32-bit mode on a file system:

```
ibrix_fs -w -f FSNAME
```

Assign or delete a standby server for the specified segments

Identify a file serving node as the standby for the specified segments:

```
ibrix_fs -b -f FSNAME -s LVLIST -h HOSTNAME
```

The command assigns *HOSTNAME* as the standby for the segments in *LVLIST* on file system *FSNAME*. To identify a standby server for all segments on a file serving node, use the following command, where *HOSTNAME2* is the standby for *HOSTNAME1*.

```
<installdirectory>/bin/ibrix_server -b -h HOSTNAME1,HOSTNAME2
```

Delete the standby assignment for the specified segments:

```
ibrix_fs -b -U -f FSNAME -s LVLIST
```

The command removes the standby assignment for the segments in *LVLIST* on file system *FSNAME*.

Options

Option	Description
-B	Marks bad segments.
-C	Enables export control.
-D	Disables quotas or Export Control.
-E	When used with -q, enables quotas on the named file system. When used with -C, enables export control on a named file system.
-F	Skips source host updates during segment migration.
-H <i>HOSTLIST</i>	A list of one or more host names.
-M	Forces migration.
-N	Skips host health checks during segment migrations.
-S <i>LV:HOSTNAME</i>	Identifies a paired logical volume (segment) name and host name. Multiple pairs are separated by commas.
-U	Unassigns backup hosts from a list of logical volumes.
-a	Creates a mountpoint with the same name as the file system and mounts the file system.
-b	Identifies a backup (standby) file serving node.

Option	Description
-c	Creates a new file system.
-d	Deletes the named file system or systems.
-e	Extends the named file system.
-f <i>FSNAME</i> or <i>FSLIST</i>	A file system or a list of file systems.
-h <i>HOSTNAME</i>	A host name.
-i	Reports on one or more file systems, and optionally on their segments.
-l	Lists file-system details.
-m	Migrates a file system or file-system segments to a new host.
-o <i>OPTION=VALUE</i>	Specifies file-system parameters; written as <i>option=value</i> pairs. Pairs valid in the current version of X9000 Software are: compat=yes no. Declares whether this file system is 32-bit compatible. maxsegs= <i>n</i> . Defines the number of segments per file system in 32-bit file systems. <i>n</i> can be set to 15, 31, 63, 127, or 255.
-p <i>PVLIST</i>	A list of one or more physical volume names, expressed as a comma-separated list (for example, d1,d2,d3) or a range (for example, d[1-3]).
-q	Enables quotas on a file system.
-s <i>LVLIST</i>	A list of one or more logical volume (segment) names, expressed as a comma-separated list (for example, ilv1,ilv2,ilv3) or a range list (for example, ilv[1-3]).
-t <i>TIERNAME</i>	Specifies the tier to which segments are assigned when creating or expanding a file system.
-w	Disables 32-bit mode.
-x	When used with -i, omits segment details from the output report.
-?	Shows the usage for this command.

See also

ibrix_fs_tune

ibrix_fs_ops

Manages directory tree quota entries.

Description

Directory tree quotas apply to a file system located at a specific directory tree. Use the `ibrix_fs_ops` command to create, delete, or view directory tree quotas. After creating a directory tree quota, use `ibrix_edquota` to set limits for the quota, including the amount of storage and the number of files that can be created.

Synopsis

Create a directory tree quota:

```
ibrix_fs_ops -D -c -f FSNAME -p PATH -n NAME
```

The `-f FSNAME` option specifies the name of the file system.

The `-p PATH` option specifies the pathname of the directory tree. If the pathname includes a space, enclose the portion of the pathname that includes the space in single quotation marks, and enclose the entire pathname in double quotation marks. For example:

```
ibrix_fs_ops -D -c -f fs48 -p "/fs48/data/'QUOTA 4'" -n QUOTA_4
```

The `-n NAME` option specifies a unique name for the directory quota. The name cannot contain a comma (,) character.

Use `ibrix_edquota` to assign usage limits to the directory tree quota.

List the directory tree quota applied to a specific file system:

```
ibrix_fs_ops -D -l -f FSNAME
```

To see the usage limits assigned to directory tree quotas, use `ibrix_edquota`.

Delete the specified directory tree quota entry and limits:

```
ibrix_fs_ops -D -d -f FSNAME -n NAME
```

The `-n NAME` option specifies the name of the directory tree quota. When the `ibrix_fs_ops` command has finished executing, issue the following command to remove the quota account for the directory tree:

```
ibrix_online_quotacheck -t -T 0 {path}
```

Options

Option	Description
-D	Indicates that this is a directory tree quota operation.
-c	Creates a directory tree quota.
-d	Deletes a directory tree quota.
-f <i>FSNAME</i>	The name of the file system.
-l	Lists directory tree quotas.
-n <i>NAME</i>	The name of the directory tree quota.
-p <i>PATH</i>	The path to the directory tree.
-?	Shows the usage for this command.

See also

`ibrix_edquota`

ibrix_fs_tune

Overrides the default file allocation behavior on file serving nodes and clients.

Description

X9000 Software allocates new files and directories to segments according to allocation policies and segment preference settings that are in effect on the host that creates the file or directory. An allocation policy controls segment selection. A segment preference setting overrides allocation policies to prefer particular segments for storage.

By default, all file serving nodes and clients use the RANDOM allocation policy and prefer all segments for storing files. These defaults are well-suited to most needs. If an application requires more sophisticated file allocation behavior, use `ibrix_fs_tune` to override the defaults as follows:

- Change the default file and directory allocation policy (the available allocation policies are listed in “File allocation policies” (page 55)).
- Prefer a pool of segments for storing all new files and directories.
- Declare the first segment in which a file or directory will be stored.
- Tune the allocation policy storage settings.



CAUTION: Changing the file allocation defaults will alter file-system storage behavior. Contact HP technical support before changing any file allocation defaults.

NOTE: If your file and directory allocation policies are different and you want to make them the same, first set a new file allocation policy, and then set the directory allocation policy to NONE. The directory allocation policy assumes the value that is set for the file allocation policy. This is the default directory allocation policy in effect until it is explicitly changed.

Because changes are specific to a given file system on a given host, different allocation policies and preferences can be set for different file systems. You can also set different allocation policies for files and directories, which can be useful for setting up data protection.

Allocation policies cannot be set on NFS and CIFS clients directly. They adhere to the policies set on the NFS/CIFS servers to which they connect.

When a client creates a new file or directory, X9000 Software chooses a segment for storing it according to the rule in effect for that node:

- Default allocation policy and segment preferences: X9000 Software applies the RANDOM primary policy and chooses a segment from among all segments.
- For non-default allocation policy and default segment preferences, X9000 Software applies the non-default allocation policy and chooses a segment from among all available segments.
- For non-default allocation policy and non-default segment preferences, X9000 Software chooses the segment according to these precedence rules:
 - a. From the pool of preferred segments, select a segment according to the allocation policy set for the client and store the file/directory there if space is available.
 - b. If all segments in the pool are full, use the AUTOMATIC algorithm to choose a segment with enough space available.

Allocation policy changes are immediately executed on file serving nodes. For X9000 clients, a policy intention is stored in the management console. When X9000 Software services start on a client, the client queries the management console for the allocation policies that it should use and then implements them. If X9000 Software services are already running on X9000 clients, you can force clients to query the management console in one of the following ways: Execute

`ibrix_client` on the X9000 clients, reboot the X9000 clients, or execute `ibrix_lwhost --a` on the X9000 clients.

Table 2 File allocation policies

AUTOMATIC	The X9000 Software selects the allocation policy.
DIRECTORY	Allocates files to the segment where its parent directory sits.
LOCAL	Allocates files to the local segments for a file serving node.
RANDOM	Allocates files to a randomly chosen segment among preferred segments. This is the default policy.
ROUNDROBIN	Allocates files to preferred segments in segment order, returning to the first segment (or the designated starting segment) when a file has been allocated to the last segment.
STICKY	Allocates files to one segment until the segment's storage limit is reached, and then moves to the next segment as determined by the AUTOMATIC algorithm.
NONE	Sets directory allocation policy only. Causes the directory allocation policy to revert to its default, which is the policy set for file allocation. Use NONE only to set file and directory allocation to the same policy.
HOST_ROUNDROBIN_NB	For clusters with more than 16 file serving nodes, takes a subset of the servers to be used for file creation and rotates this subset on a regular periodic basis. This policy should be used only under the direction of HP technical support.

Synopsis

Set an allocation policy:

```
ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -p POLICY [-S STARTSEGNUM] [-R]
```

The policy is applied to the specified file serving nodes and X9000 clients (*HOSTLIST*) or hostgroups (*GROUPLIST*) for file system *FSNAME*, starting at the segment with the most available storage.

To specify the starting segment for applying the policy, include the `-S STARTSEGNUM` option.

To set the policy for files, omit the `-R` option. To set the policy for directories, include the `-R` option.

Reset the specified hosts or hostgroups to the default allocation policy for the file system:

```
ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -p -U
```

Prefer a pool of segments for the specified hosts or hostgroups:

```
ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -s LVNAMELIST
```

The segments to be included in the pool are specified by logical volume name (*LVNAMELIST*).

To prefer a pool of segments for NFS/CIFS clients, specify only NFS/CIFS servers in the *HOSTLIST*.

To prefer a pool of segments for X9000 clients, either specify X9000 clients in the *HOSTLIST*, or specify a hostgroup in the *GROUPLIST*. To prefer all X9000 clients, specify the `clients` hostgroup.

Prefer a pool of segments on specified file serving nodes or hostgroups in accordance with the `-s` option:

```
ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -s {SEGNUMLIST|ALL|LOCAL}
```

- *SEGNUMLIST*: Takes a list of segment names.
- *ALL*: A keyword that allows clients to write to any segment on the listed file serving nodes. This restores the default segment preferences for the clients.
- *LOCAL*: A keyword that allows clients to write only to the local segments on the specified file serving nodes.

To prefer a pool of segments for NFS/CIFS clients, specify only NFS/CIFS servers in the *HOSTLIST*. To prefer a pool of segments for X9000 clients, lists the X9000 clients in *HOSTLIST* or specify a hostgroup in the *GROUPLIST*. To prefer all X9000 clients, specify the *clients* hostgroup.

Prefer a tier for the specified hosts or hostgroups:

```
ibrix_fs_tune -f FSNAME {-h HOSTLIST | -g GROUPLIST} -t TIERNAME
```

Tune allocation policy storage settings globally on the specified file system:

```
ibrix_fs_tune -f FSNAME -O [-p POLICY] [-S STARTSEGNUM] [-P prealloc (KB)] [-r readahead (KB)] [-N NFS readahead (KB)]
```

The following settings can be tuned:

- Default allocation policy for all hosts
- Default starting segment number for applying settings
- Number of KB to allocate for a file at one time (*prealloc*)
- Number of KB that X9000 Software will pre-fetch (*readahead*)
- Number of KB that X9000 Software will pre-fetch under NFS (*NFS readahead*)

The default for *prealloc* is 256 KB. The defaults for *readahead* and *NFS readahead* are 128 KB.

List preferred segments or allocation policy for the specified hosts, hostgroups, or file system:

```
ibrix_fs_tune -l [-h HOSTLIST | -g GROUPLIST] [-f FSNAME]
```

The *-S* option lists preferred segments. The *-P* option lists allocation policy.

List the preferred segments on the specified file serving nodes or hostgroups:

```
ibrix_fs_tune {-h HOSTLIST | -g GROUPLIST} -l -S
```

Options

Option	Description
<i>-N NFS readahead</i>	Specifies the number of KB that X9000 Software pre-fetches under NFS. The default value is 128 KB.
<i>-O</i>	Sets one or more allocation policy tuning options.
<i>-P prealloc</i>	Sets the number of KB a file system pre-allocates to a file. Default: 256 KB.
<i>-R</i>	Sets allocation policy for directory segments only.
<i>-S SEGMENT_NUMBER</i>	Specifies a segment.
<i>-S SEGNUMLIST ALL LOCAL</i>	Specifies a list of segments to prefer, prefers all segments, or restricts clients to using segments local to the file serving nodes listed in the command. <i>ALL</i> and <i>LOCAL</i> are keywords.
<i>-S STARTSEGNUM</i>	Identifies the first segment to which an allocation policy is applied in a file system.
<i>-U</i>	Resets an allocation policy.
<i>-f FSNAME</i>	A file system.
<i>-g GROUPLIST</i>	A list of hostgroups.
<i>-h HOSTLIST</i>	A list of one or more file serving nodes or X9000 clients
<i>-l u g f</i>	Lists allocation policies for users, groups, or file systems.
<i>-p POLICY</i>	Sets allocation policy on a file system and optionally on a list of hosts or groups. Policy type names are case-sensitive and must be entered as shown: <i>ROUNDROBIN</i> , <i>STICKY</i> , <i>DIRECTORY</i> , <i>LOCAL</i> , <i>RANDOM</i> , <i>NONE</i> , <i>HOST_ROUNDROBIN_NB</i> .

Option	Description
-r readahead	Specifies the number of KB that X9000 Software will pre-fetch. The default value is 128 KB.
-s <i>LVNAME</i> or <i>LVNAMELIST</i>	A list of one or more logical volume names (segments).
-t <i>TIERNAME</i>	A tier on a file system.
-?	Shows the usage for this command.

See also

ibrix_fs, ibrix_hostgroup

ibrix_fsck

Analyzes inconsistencies in a file system.

Description

If a file system shows evidence of inconsistencies, contact HP technical support. You may be instructed to run `ibrix_fsck` with corrective options that are not described in this guide.

⚠ CAUTION: Do not run `e2fsck` or any other off-the-shelf `fsck` program on any part of a file system. Doing so will damage the file system!

`ibrix_fsck` checks segments for file-system inconsistencies and optionally repairs them. The command runs in four phases and has two running modes, analytical and corrective. Run all the phases in order, without skipping any.

- Phase 0 checks host connectivity and the consistency of the segment byte blocks and repairs them in corrective mode.
- Phase 1 checks segments, repairs them in corrective mode, and stores the results locally.
- Phase 2 checks the file system, repairs it in corrective mode, and stores the results locally.
- Phase 3 moves files from `lost+found` on each segment to the global `lost+found` directory on the file system's root segment.

NOTE: During an `ibrix_fsck` run, an `INFSCCK` flag is set on the file system to protect it. If an error occurs during the run, you must explicitly clear the `INFSCCK` flag with `ibrix_fsck -f FSNAME -C` or you will be unable to mount the file system.

Observe the following requirements when executing `ibrix_fsck`:

- Unmount the file system for phases 0 and 1 and mount the file system for phases 2 and 3.
- Turn off automated failover by executing `ibrix_host -m -U -h SERVERNAME`.
- Unmount all NFS clients and stop NFS on the servers.
- If running X9000 Software HA, put it into administrative mode (`ibrix_fmha -a`) for the duration of the analysis.

Synopsis

Clear flags for the specified file system:

```
ibrix_fsck -f FSNAME -C
```

Runs phase 0 in analytic mode:

```
ibrix_fsck -p 0 -f FSNAME [-s LVNAME] [-c]
```

The command can be run on the specified file system or optionally only on the specified segment `LVNAME`.

Run phase 1 in analytic mode:

```
ibrix_fsck -p 1 -f FSNAME [-s LVNAME] [-c] [-B BLOCKSIZE] [-b  
ALTSUPERBLOCK]
```

The command can be run on file system `FSNAME` or optionally only on segment `LVNAME`. This phase can be run with a specified block size and an alternate superblock number. For example:

```
ibrix_fsck -p 1 -f ifs1 -B 4096 -b 12250
```

NOTE: If phase 1 is run in analytic mode on a mounted file system, false errors can be reported.

Run phase 2:

```
ibrix_fsck -p 2 -f FSNAME [-s LVNAME] [-c] [-o "options"]
```

The command can be run on the specified file system or optionally only on segment *LVNAME*. Use *-o* to specify any options.

Run phase 3:

```
ibrix_fsck -p 3 -f FSNAME [-c]
```

Clear the INFCK flag:

```
ibrix_fsck -f FSNAME [-C]
```

Options

Option	Description
<i>-B BLOCKSIZE</i>	Specifies a block size in KB.
<i>-C</i>	Clears the INFCK flag.
<i>-b ALTSUPERBLOCK</i>	Specifies the number of an alternate superblock, generally used only if the normal superblock is corrupted.
<i>-c</i>	Runs <i>ibrix_fsck</i> in corrective mode (the default is analytical mode).
<i>-f FSNAME</i>	A file system.
<i>-o "options"</i>	Options for the command.
<i>-p 0 1 2 3</i>	Runs the specified phase (0, 1, 2, or 3).
<i>-s LVNAME</i>	Specifies a logical volume (segment) name.
<i>-?</i>	Shows the usage for this command.

See also

ibrix_host, *ibrix_umount*

ibrix_ftpconfig

Manages configuration profiles for the FTP service.

Description

A configuration profile specifies a set of global FTP parameters that are in effect on the file serving nodes listed in the profile. The `vsftpd` service starts on these nodes when the cluster services are started. Only one configuration profile can be in effect on a particular node.

When you create an FTP share, you will need to associate it with a configuration profile. The FTP settings assigned to the share will be applied to the file serving nodes specified in the configuration profile.

Synopsis

Add a configuration profile:

```
ibrix_ftpconfig -a PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"passive_enable=true,maxclients=200,..."`. To see a list of available settings for the share, use `ibrix_ftpconfig -L`.

Modify a configuration profile:

```
ibrix_ftpconfig -m PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

Delete a configuration profile:

```
ibrix_ftpconfig -d PROFILENAME
```

Display current information about a configuration profile:

```
ibrix_ftpconfig -i PROFILENAME [-v level]
```

Use `-v 1` to display detailed information.

Display current information about configuration profiles on the specified hosts:

```
ibrix_ftpconfig -i -h HOSTLIST [-v level]
```

Use `-v 1` to display detailed information.

List all configuration profiles:

```
ibrix_ftpconfig -l [-v level]
```

Use `-v 1` to display detailed information.

List the default settings for a configuration profile:

```
ibrix_ftpconfig -L
```

Options

Option	Description
<code>-L</code>	Lists default profile settings.
<code>-S SETTINGLIST</code>	Specifies the settings that are to be applied to a configuration profile. Use a comma to separate the settings, and enclose the settings in quotation marks, such as <code>"passive_enable=true,maxclients=200,..."</code> . To see a list of available settings for the share, use <code>ibrix_ftpconfig -L</code> .
<code>-a</code>	Adds a configuration profile.
<code>-d</code>	Deletes a configuration profile.
<code>-h HOSTLIST</code>	Specifies one or more file serving nodes.

Option	Description
-i	Displays current information for the specified configuration profile, or displays current profile information for specified file serving nodes.
-l	Lists all configuration profiles.
-m	Modifies the specified configuration profile.
-v <i>level</i>	Displays detailed information. Enter 1 as the <i>level</i> .
-?	Shows the usage for this command.

See also

`ibrix_ftpshare`

ibrix_ftpshare

Manages FTP shares.

Description

An FTP share provides access to data on X9000 Software file systems. FTP clients access the shares using standard FTP and FTPS protocol services.

The `ibrix_ftpshare` command can be used to add a share to an existing file system. You can create multiple shares having the same physical path, but with different sets of properties, and then assign users to the share that is appropriate for their usage.

NOTE: The file system must be mounted when you add the share.

Synopsis

Add an FTP share:

```
ibrix_ftpshare -a SHARENAME -c PROFILENAME -f FSNAME -p dirpath -I  
IP-address:Port [-u USERLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"browseable=true,readonly=true,..."`. The settings should include the IP addresses or Virtual IP addresses that clients will use to access the share. Enter the addresses as IP address:port pairs. To see a list of the allowable settings for the share, use `ibrix_ftpshare -L`.

For the `-I` option, use a semicolon to separate the IP address:port settings and enclose the settings in quotation marks, such as `"ip1:port1;ip2:port2;..."`.

The `-U` option allows you to specify the allowed permissions and an optional custom directory path for specific users. Use commas to separate the users (for example: `user1=access[;custom_user_path],USER2=access[;custom_user_path],...`). In the command, access can be `r/w/rw`.

Modify an FTP share:

```
ibrix_ftpshare -m SHARENAME -c PROFILENAME [-f FSNAME -p dirpath] -I  
IP-address:Port [-u USERLIST] [-S SETTINGLIST]
```

Delete an FTP share:

```
ibrix_ftpshare -d SHARENAME -c PROFILENAME
```

Delete all FTP shares associated with a file system:

```
ibrix_ftpshare -d -f FSNAME
```

Display information about an FTP share:

```
ibrix_ftpshare -i SHARENAME -c PROFILENAME [-v level]
```

Use `-v 1` to display detailed information.

List all FTP shares, or shares associated with a specific profile:

```
ibrix_ftpshare -l -c PROFILENAME [-v level]
```

Use `-v 1` to display detailed information.

List FTP shares associated with a specific file system:

```
ibrix_ftpshare -l -f FSNAME [-v level]
```

Use `-v 1` to display detailed information.

List the valid settings for a share:

```
ibrix_ftpshare -L
```

Options

Option	Description
-I <i>IPaddrs:Port</i>	Specifies the IP addresses/ports assigned to this share.
-L	Lists valid FTP share settings.
-S <i>SETTINGLIST</i>	Specifies the settings to apply to an FTP share. Use a comma to separate the settings, and enclose the settings in quotation marks, such as "browseable=true,readonly=true,...". To see a list of available settings for the share, use <code>ibrix_ftpshare -L</code> .
-a <i>SHARENAME</i>	Adds an FTP share.
-c <i>PROFILENAME</i>	Specifies a configuration profile.
-d	Deletes an FTP share.
-f <i>FSNAME</i>	Specifies a file system.
-i	Displays information about FTP shares.
-l	Lists FTP shares.
-m <i>SHARENAME</i>	Modifies the specified FTP share.
-p <i>dirpath</i>	Specifies the path to the share.
-u <i>USERLIST</i>	Specifies users and their access permissions. Use commas to separate the users (for example: <code>user1=access[;custom_user_path],USER2=access[;custom_user_path],...</code>). The value of <code>access</code> can be <code>r/w/rw</code> .
-v <i>level</i>	Displays detailed information. Enter 1 as the <i>level</i> .
-?	Shows the usage for this command.

See also

`ibrix_ftpconfig`

ibrix_haconfig

Reports whether X9000 Software high availability features are configured for file serving nodes.

Description

`ibrix_haconfig` tests whether the following high availability features are configured for file serving nodes and provides either a summary or a detailed report of the results:

- Programmable power source
- Standby server or standby segments
- Cluster and user network interface monitors
- Standby network interface for each user network interface
- HBA port monitoring
- Whether automated failover is turned on

The summary report returns a result status for each high availability feature, for each tested file serving node, and optionally for their standbys:

- **Passed:** The feature is configured.
- **Warned:** The feature is not configured, but the result is ambiguous. For example, the configuration checker cannot know whether the absence of discovered HBAs indicates that the HBA monitoring feature was not set up or that HBAs are not physically present on the tested servers.
- **Failed:** The feature is not configured.

The report includes an overall result status for all tested file serving nodes and describes details about the checks performed on each high availability feature. By default, the report includes details only about checks that received a **Failed** or a **Warned** result. You can expand the report to include details about checks that received a **Passed** result.

Synopsis

Determine whether X9000 Software high availability is configured and report overall results:

```
ibrix_haconfig -l [-h HOSTLIST] [-f] [-b]
```

The command checks either all file serving nodes or only those nodes in *HOSTLIST*, and reports the overall results for all tested hosts. To include standbys for file serving nodes in the check, include the `-b` option. To list only the file serving nodes that failed a health test, include the `-f` option.

Determine whether X9000 Software high availability is configured and report detailed failed/warned results:

```
ibrix_haconfig -i [-h HOSTLIST] [-f] [-b] [-v]
```

The command checks either all file serving nodes or only those nodes in *HOSTLIST* and reports the overall results for all tested hosts. Detailed information is provided for configuration checks that received a **Failed** or **Warned** result.

To include standbys for file serving nodes in the check, include the `-b` option. To include information only about file serving nodes that failed a health test, include the `-f` option. To expand the report to include detailed information about configuration checks that received a **Passed** result, include the `-v` option.

Options

Option	Description
-b	Includes standby servers in the configuration check.
-f	Reports only on failed file serving nodes.
-h <i>HOSTLIST</i>	A list of file serving nodes.
-i	Displays a detailed high availability configuration report.
-l	Displays a summary high availability configuration report.
-v	Displays details about the configuration checks that received a <code>Passed</code> result.
-?	Shows the usage for this command.

ibrix_hba

Discovers HBAs, identifies standby pairings of HBA ports, monitors HBA ports for failure, deletes HBAs from the configuration database, and lists information about HBAs.

Description

`ibrix_hba` sets up HBA monitoring. When a monitored HBA port fails, the management console fails the host over to its standby if a standby has been declared and if automated failover is turned on.

The `-b` option identifies two multipath HBAs that have been set to operate as a failover pair. Built-in standby switching is vendor-supplied and is not present in all HBA models. If the operating HBA fails, the standby HBA is activated. If both HBAs fail and HBA monitoring is enabled, the management console fails over the host.

In some of these HBA commands, worldwide names (WWN) are assigned as parameter values. These are either worldwide node names (WWNN) or worldwide port names (WWPN). The WWPN is the name an HBA presents when logging into a SAN fabric. Worldwide names consist of 16 hexadecimal digits grouped in pairs. In X9000 Software, these are written as dot-separated pairs such as 21.00.00.e0.8b.05.05.04.

Synopsis

Discover HBA ports and add them to the configuration database:

```
ibrix_hba -a [-h HOSTLIST]
```

Use the `-h HOSTLIST` option to limit the command to specific file serving nodes.

Enable HBA port monitoring on a file serving node:

```
ibrix_hba -m -h HOSTNAME -p <PORT>
```

Disable HBA port monitoring on a file serving node:

```
ibrix_hba -m -U -h HOSTNAME -p PORT
```

Add an HBA standby pair:

```
ibrix_hba -b -P WWPN1:WWPN2 -h HOSTNAME
```

The command identifies `WWPN1` and `WWPN2` as a standby pair on file serving node `HOSTNAME`.

Delete an HBA standby pair:

```
ibrix_hba -b -U -P WWPN1:WWPN2 -h HOSTNAME
```

The command deletes the standby pairing of ports `WWPN1` and `WWPN2` on file serving node `HOSTNAME` from the configuration database.

Delete the HBA entries in `WWNNLIST` from the configuration database for a file serving node:

```
ibrix_hba -d -h HOSTNAME -w WWNNLIST
```

List information about HBA ports:

```
ibrix_hba -l [-h HOSTLIST]
```

Use the `-h HOSTLIST` option to limit the command to specific file serving nodes.

Options

Option	Description
<code>-P WWPN</code>	An HBA's worldwide port name, specified as 16 hex digits grouped into eight pairs separated by dots (for example, 20.00.12.34.56.78.9a.bc).
<code>-U</code>	Stops HBA monitoring or deletes a standby pairing from the configuration database.

Option	Description
-a	Adds an HBA port definition.
-b	Identifies two ports as a standby pair.
-d	Deletes an HBA port definition.
-h <i>HOSTNAME</i> or <i>HOSTLIST</i>	One or more file serving nodes.
-l	For each host, displays its WWNN, WWPN, backup WWPN, and whether HBA monitoring is turned on.
-m	Turns on HBA monitoring.
-p <i>PORT</i>	Specifies the port number used for HBA monitoring.
-w <i>WWNN</i> or <i>WWNNLIST</i>	An HBA's worldwide node name or a list of such names, specified as 16 hex digits grouped into eight dot-separated pairs (for example, 20.00.12.34.56.78.9a.bc). If you specify multiple names, use commas to separate the names (for example, 50.01.43.80.02.51.bf.c7,50.01.43.80.02.51.bf.c5).
-?	Shows the usage for this command.

ibrix_health

Checks the functional health of file serving nodes and X9000 clients.

Description

`ibrix_health` runs a set of health checks on all hosts or the specified hosts and reports the results. To constrain the report to show only hosts that failed, include the `-f` option.

On both file serving nodes and X9000 clients, the check determines whether information maps on the tested hosts are consistent with the configuration database.

Checks run on file serving nodes do the following:

- Ping remote file serving nodes that share a network with the test hosts but only if the test hosts do not have an open connection to the remote servers. Remote servers that can be pinged may not be connected to a test host because of a Linux or X9000 Software problem. Remote servers that cannot be pinged may be down or have a network problem.
- If test hosts are assigned to be network interface monitors, ping their monitored interfaces to assess the health of the connection.
- Determine whether hosts can read their physical volumes.

Health checks are reported as Passed, Failed, or Warning:

- **Passed:** All tested hosts and standbys passed every health check.
- **Failed:** One or more tested hosts failed a health check. The health status of standbys for file serving nodes is not included when this result is calculated.
- **Warning:** A suboptimal condition that may require attention was found on one or more tested hosts or standbys.

The terms have the same meaning for individual file serving nodes and X9000 clients. A standby for a file serving node is not considered when the result is determined for the node.

The basic (`-l`) report includes the following for each host: hostname, test result (passed/failed), host type (server/client), state (up/down), and date of last update.

The detailed report consists of a summary report plus the following data: a summary of the test results; host information such as operational state, performance data, and version data; non-default host tunings; and the results of the health checks. By default, the "Result Information" field in a detailed report provides data only for health checks that received a Failed or a Warned result. You can expand a detailed report to also provide data about checks that received a Passed result, as well as details about the file system and segments.

Synopsis

Perform health checks and provide a summary report of the results:

```
ibrix_health -l [-h HOSTLIST] [-f] [-u] [-b] [-t SCRIPTNAMELIST]
```

The command is run on the specified file serving nodes and X9000 clients. The options are:

- `-h HOSTLIST:` Runs health checks on the specified file serving nodes and X9000 clients.
- `-f:` Reports only hosts that failed a test.
- `-u:` Reports only hosts known to be up.
- `-b:` Includes standbys for file serving nodes.
- `-t SCRIPTNAMELIST:` Runs the specified scripts and include them in the health report.

Display a summary health report:

```
ibrix_health -i -h HOSTLIST [-f] [-b] [-s] [-v] [-t SCRIPTNAMELIST]
```

The options are:

- `-h HOSTLIST:` Runs health checks on the specified file serving nodes and X9000 clients.

- f: Reports only hosts that failed a test.
- b: Includes standbys for file serving nodes.
- s: Includes details about the file system and segments.
- v: Includes details about checks that received a Passed result.
- t *SCRIPTNAMELIST*: Includes details about the specified scripts.

Options

Option	Description
-b	Includes standbys for file serving nodes in the check.
-f	Includes only failed hosts in the report.
-h <i>HOSTLIST</i>	A list of one or more hosts.
-i	Displays a detailed health report.
-l	Displays a summary health report.
-s	Includes file system and remote host information in the health report.
-t <i>SCRIPTNAMELIST</i>	Runs the specified scripts and includes them in the health report.
-u	Includes only hosts known to be up in the report.
-v	Lists details about the test that was performed.
-?	Shows the usage for this command.

ibrix_healthconfig

Manages health check sets.

Description

Use this command to create or update sets of health checks. You can also delete health check sets display existing health check sets.

Synopsis

Create a set of health checks:

```
ibrix_healthconfig -c -n SETNAME -t HEALTH_CHECKLIST
```

SETNAME is the name for the set being created and *HEALTH_CHECKLIST* lists the checks to be added.

Add health checks to health check sets:

```
ibrix_healthconfig -u -A -n SETLIST -t HEALTH_CHECKLIST
```

Remove health checks from health check sets:

```
ibrix_healthconfig -u -D -n SETLIST -t HEALTH_CHECKLIST
```

Delete health sets:

```
ibrix_healthconfig -d -n SETLIST
```

Display detailed information about health check sets:

```
ibrix_healthconfig -i -n SETLIST
```

Lists the available health checks:

```
ibrix_healthconfig -L
```

Options

Option	Description
-A	Adds health checks to a set.
-D	Removes health checks from a set.
-L	Lists the names of available health checks.
-c	Creates a set of health checks.
-d	Deletes a set of health checks.
-i	Displays detailed information about health check sets.
-n <i>SETLIST</i>	The name of a health check set.
-t <i>HEALTH_CHECKLIST</i>	List of health checks.
-u	Updates a set of health checks.
?	Shows the usage for this command.

ibrix_host_tune

Sets and lists tuning parameters for file serving nodes and X9000 clients.

Description

- △ **CAUTION:** The default values for the host tuning parameters are suitable for most cluster environments. Because changing parameter values can alter file-system performance, HP recommends that you exercise caution before implementing any changes, or do so only under the guidance of HP technical support.

Host tunings are immediately executed for file serving nodes. For X9000 clients a tuning intention is stored in the management console. When HP services start on a X9000 client, the client obtains its settings from the management console and implements them. If HP services are already running on X9000 clients when new parameter values are set, you can force the clients to query the management console for the settings by running `ibrix_client` or `ibrix_lwhost --a` on the X9000 clients, or by rebooting them.

To locally override host tunings on a single X9000 client, run `ibrix_lwhost`. All of the management console commands for tuning hosts include an `-h HOSTLIST` option that can take one or more hostgroups. Setting host tunings on a hostgroup is a convenient way to tune a set of X9000 clients all at once. To set the same host tunings on all X9000 clients, specify the built-in `clients` hostgroup.

Synopsis

Set tunings on file serving nodes and X9000 clients:

```
ibrix_host_tune -S {-h HOSTLIST|-g GROUPLIST} -o OPTIONLIST
```

The command sets host tunings (*OPTIONLIST*) on specified file serving nodes and X9000 clients (*HOSTLIST*), or on specified hostgroups (*GROUPLIST*). For example, to set host tunings on all X9000 clients, use `-g clients`.

Reset host tunings to their default values:

```
ibrix_host_tune -U {-h HOSTLIST|-g GROUPLIST} -n OPTIONS
```

The host tunings (*OPTIONS*) are reset on the specified file serving nodes and X9000 clients (*HOSTLIST*), or on specified hostgroups (*GROUPLIST*). For example, to set host tunings on all X9000 clients, use `-g clients`. The values that are restored depend on what is supplied to the *HOSTLIST* option:

- File serving nodes: The default host tunings are restored.
- Hostgroups: The host tunings that are in effect for the parent of the specified hostgroups are restored.
- X9000 clients: The host tunings that are in effect for the default `clients` hostgroup are restored.

Lists host tunings for the specified file serving nodes, X9000 clients, and hostgroups:

```
ibrix_host_tune -l [-h HOSTLIST] [-n OPTIONS]
```

Only settings that differ from the default values are listed. To list settings for specific parameters, include the `-n OPTIONS` option.

To list settings for all file serving nodes and X9000 clients, omit the `-h HOSTLIST` option.

List host tunings for the specified groups:

```
ibrix_host_tune -l {-g GROUPLIST | -G } [-n OPTIONS]
```

Only settings that differ from the default values are listed. To list settings for specific parameters, include the `-n OPTIONS` option.

List default values and the range of values for host-tuning parameters:

```
ibrix_host_tune -L
```

Set host configuration options for the specified hosts:

```
ibrix_host_tune -C [-h HOSTLIST] -o OPTIONLIST
```

OPTIONS is a list of configuration options separated by commas.

Reset host configuration options to the default values:

```
ibrix_host_tune -R [-h HOSTLIST] -n OPTIONS
```

OPTIONS is a list of configuration options separated by commas.

List non-default configuration settings for the specified hosts:

```
ibrix_host_tune -q [-h HOSTLIST] [-n OPTIONS]
```

OPTIONS is a list of configuration options separated by commas.

List default configuration options:

```
ibrix_host_tune -Q
```

Set the communications protocol on file serving nodes, X9000 clients, or hostgroups:

```
ibrix_host_tune -p {UDP|TCP} {-h HOSTLIST|-g GROUPLIST}
```

The default protocol is TCP. To set the protocol on all X9000 clients, use *-g clients*.

Set server threads on file serving nodes, X9000 clients, or hostgroups:

```
ibrix_host_tune -t THREADCOUNT {-h HOSTLIST|-g GROUPLIST}
```

To set server threads on all X9000 clients, use *-g clients*.

Set admin threads on file serving nodes, X9000 clients, or hostgroups:

```
ibrix_host_tune -a THREADCOUNT {-h HOSTLIST|-g GROUPLIST}
```

To set admin threads on all X9000 clients, use *-g clients*.

Options

Option	Description
-C	Sets configuration tuning options.
-G	Includes all groups
-L	Lists default tuning options.
-Q	Lists default configuration tuning values.
-R	Resets configuration tuning options.
-S	Sets host tuning options.
-U	Resets host tuning options to default values.
-a <i>THREADCOUNT</i>	Sets the number of admin threads. The default value is 3.
-g <i>GROUPLIST</i>	Identifies one or more hostgroups.
-h <i>HOSTLIST</i>	Identifies one or more file serving nodes or X9000 clients.
-l	Lists host tuning options.
-n <i>OPTIONS</i>	Specifies a list of host tuning options, separated by commas, and in the format <i>option1,option2...</i>
-o <i>OPTIONLIST</i>	Specifies a list of comma-separated option/value pairs, separated by commas, and in the format <i>option1=value1,option2=value2...</i>

Option	Description
-p tcp udp	Sets a communications protocol (UDP or TCP).
-q	Queries configuration tuning parameter values.
-t <i>THREADCOUNT</i>	Sets the number of server threads. The default value is 10.
-?	Shows the usage for this command.

See also

ibrix_hostgroup, ibrix_lwhost

ibrix_hostgroup

Creates and manages hostgroups.

Description

A hostgroup is a named set of X9000 clients. Hostgroups provide a convenient way to centrally manage clients. Group-wide procedures include:

- Mounting file systems.
- Preferring or unpreferring a network interface for a hostgroup. (You can locally override the preference on individual X9000 clients with `ibrix_lwhost`.)
- Tuning host parameters.
- Setting allocation policies.

When `ibrix_mount`, `ibrix_host_tune`, or `ibrix_fs_tune` is executed on a hostgroup, the command is stored on the management console for each X9000 client until the next time a X9000 client contacts it. This happens when the X9000 client reboots, when X9000 Software services start on the X9000 client, or when it is forced to query the management console. When contacted, the management console replies to the X9000 clients with information about commands that have been executed on hostgroups to which they belong. The X9000 clients then use this information to mount file systems or implement host tunings or allocation policy settings.

To force a query on a X9000 client, execute `ibrix_lwmount -a` if you executed `ibrix_mount` or `ibrix_fs_tune` on the hostgroup, or run `ibrix_lwhost --a` if you executed `ibrix_host_tune`.

All X9000 clients belong to the default `clients` hostgroup, which can be used to globally mount file systems, tune hosts, or set allocation policy. To perform different actions on different sets of X9000 clients, create additional hostgroups and assign X9000 clients to them. This involves creating a hostgroup tree whose root element is the `clients` hostgroup. Each hostgroup in a tree can have exactly one parent, and a parent can have multiple children. To create the first level of nodes beneath the root, execute `ibrix_hostgroup -c -g GROUPNAME`. To create subsequent levels of nodes, execute `ibrix_hostgroup -c -g GROUPNAME -p PARENT`.

After creating a hostgroup, assign X9000 clients to it. You can do this manually or have the management console automatically assign the clients when they are registered. To set up automatic hostgroup assignments, define a domain rule for hostgroups. A domain rule restricts hostgroup membership to X9000 clients on a particular cluster subnet. The management console uses the IP address that you specify for X9000 clients when you register them to perform a subnet match and sort the X9000 clients into hostgroups based on domain rules. Setting domain rules on hostgroups thus provides a convenient way to centrally manage mounting, tuning, and allocation policies on different subnets of X9000 clients.

In a hostgroups tree, operations performed on lower-level nodes take precedence over operations performed on higher-level nodes. This means that you can effectively establish global X9000 client mounts, host tunings, and allocation policies and then override them for specific hostgroups.

Synopsis

Create a hostgroup tree:

```
ibrix_hostgroup -c -g GROUPNAME [-D DOMAIN] [-p PARENT]
```

To create one level of hostgroups beneath the root `clients` hostgroup, execute the command once for every hostgroup that you want to add, but omit the `-p PARENT` option.

To create additional levels of hostgroups, execute the command for each hostgroup to be added and include the `-p PARENT` option.

To set a domain rule on a hostgroup, include the `-D DOMAIN` option.

Add a domain rule to a hostgroup:

```
ibrix_hostgroup -a -g GROUPNAME -D DOMAIN
```

Delete one or more hostgroups from a hostgroup tree:

```
ibrix_hostgroup -d -g GROUPLIST
```

Prefer a network interface for communication from a hostgroup to a destination host:

```
ibrix_hostgroup -n -g HOSTGROUP -A DESTHOST/IFNAME
```

To prefer a network interface for all X9000 clients, use `-g clients`.

Use the cluster interface for communications from a hostgroup to a destination host:

```
ibrix_hostgroup -n -g HOSTGROUP -D DESTHOST
```

To unprefer a network interface for all X9000 clients, use `-g clients`.

Migrate an X9000 client to a different hostgroup:

```
ibrix_hostgroup -m -g GROUP -h MEMBER
```

Display detailed information for one or more hostgroups:

```
ibrix_hostgroup -i [-g GROUPLIST]
```

The information includes, as applicable, the hostgroup name, parent, domain, protocol, number of admin threads, and number of server threads.

List all hostgroups by name, parent, and domain:

```
ibrix_hostgroup -l [-g GROUPLIST]
```

To list specific hostgroups, include the `-g GROUPLIST` option.

Options

Option	Description
<code>-A DESTHOST/ IFNAME</code>	Adds a preferred NIC (<i>IFNAME</i>) to use when communicating with this destination host.
<code>-D DOMAIN</code> or <code>DESTHOST</code>	Specifies a domain rule, an IP address that corresponds to a user interface, or a destination host.
<code>-a</code>	Changes a domain attribute on a hostgroup.
<code>-c</code>	Creates a hostgroup.
<code>-d</code>	Deletes a hostgroup from the hostgroup tree.
<code>-g GROUPNAME</code> or <code>GROUPLIST</code>	Identifies the hostgroup or list of hostgroups that the command is acting on.
<code>-h MEMBER</code>	Specifies the name of a host to add to a hostgroup.
<code>-i</code>	Displays detailed information on all hostgroups or a named hostgroup.
<code>-l</code>	Lists parent and domain for all hostgroups.
<code>-m</code>	Moves a host to another hostgroup.
<code>-n</code>	Prefers or unprefers a network interface.
<code>-p PARENT</code>	Specifies a parent node in a hostgroup tree.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_host`, `ibrix_host_tune`, `ibrix_fs_tune`, `ibrix_lwhost`, `ibrix_lwmount`, `ibrix_mount`

ibrix_hostpower

Sets up a file serving node power source for failover and lists information about the power sources.

Description

Use `ibrix_hostpower` to:

- Associate a file serving node with an APC power source slot. (It is not necessary to specify a slot for an integrated power source because those power sources are always associated to slot 1.)
- Change the APC slot association for a file serving node.
- Dissociate a file serving node from a power source.
- List information about power sources.

Before associating file serving nodes with APC power sources, run `ibrix_powersrc` to identify the power sources.

Synopsis

Associate a file serving node with a slot on an APC power source:

```
ibrix_hostpower -a -i SLOT_ID -s POWERSOURCE -h HOSTNAME
```

Associate a file serving node with a different power source slot:

```
ibrix_hostpower -m -i FROM_SLOT_ID, TO_SLOT_ID -s POWERSOURCE -h HOSTNAME
```

Dissociate a file serving node from slots on a power source:

```
ibrix_hostpower -d [-s POWERSOURCE [-i SLOT]] -h HOSTNAME
```

To dissociate a host from a specific slot on an APC power source, include the `-i SLOT` option. Otherwise, the command dissociates the host from all slots.

To dissociate a file serving node from an integrated power source, omit the `-i SLOT` option. Power sources managed by these interfaces are connected to slot 1 by default.

List the power sources assigned to all file serving nodes or the specified nodes:

```
ibrix_hostpower -l [-h HOSTLIST]
```

Options

Option	Description
-a	Associates a hostname with a power source.
-d	Dissociates a hostname from a power source.
-h <i>HOSTNAME</i> or <i>HOSTLIST</i>	A file serving node name or a list of file serving node names.
-i <i>SLOT</i> or <i>SLOT_ID</i>	Identifies the number of a slot on a power source.
-l	Lists all associations between hosts and power sources.
-m	Modifies a power source slot assignment.
-s <i>POWERSOURCE</i>	Specifies the name of the power source.
-?	Shows the usage for this command.

See also

`ibrix_powersrc`

ibrix_httpconfig

Manages configuration profiles for the HTTP service.

Description

A configuration profile specifies a set of global HTTP parameters in effect on the file serving nodes listed in the profile.

When you create an HTTP share, you will need to associate it with a configuration profile. One or more HTTP virtual hosts are also associated with the profile and specify the IP addresses/ports that clients will use to access the HTTP shares.

Synopsis

Add a configuration profile:

```
ibrix_httpconfig -a PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"keepalive=true,maxclients=200,..."`. To see a list of available settings for the share, use `ibrix_httpconfig -L`.

You can specify multiple values for the `ports` and `sslports` settings. Use semicolons to separate the values, and use commas to separate the settings. For example:

```
ibrix_httpconfig -a profile1 -h node1 -S  
"ports=80;81;82,sslports=443;444,keepalive=true"
```

Modify a configuration profile:

```
ibrix_httpconfig -m PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

Delete a configuration profile:

```
ibrix_httpconfig -d PROFILENAME
```

Display current information about a configuration profile:

```
ibrix_httpconfig -i PROFILENAME [-v level]
```

Use `-v 1` to display detailed information.

Display current information about configuration profiles on the specified hosts:

```
ibrix_httpconfig -i -h HOSTLIST [-v level]
```

Use `-v 1` to display detailed information.

List all configuration profiles:

```
ibrix_httpconfig -l [-v level]
```

Use `-v 1` to display detailed information.

List the default settings for a configuration profile:

```
ibrix_httpconfig -L
```

Options

Option	Description
<code>-L</code>	Lists default profile settings.
<code>-S SETTINGLIST</code>	Specifies the settings that are to be applied to a configuration profile. Use a comma to separate the settings, and enclose the settings in quotation marks, such as <code>"keepalive=true,maxclients=200,..."</code> . To see a list of available settings for the share, use <code>ibrix_httpconfig -L</code> .
<code>-a</code>	Adds a configuration profile.

Option	Description
-d	Deletes a configuration profile.
-h <i>HOSTLIST</i>	Specifies one or more file serving nodes.
-i	Displays current information for the specified configuration profile, or displays current profile information for specified file serving nodes.
-l	Lists all configuration profiles.
-m	Modifies the specified configuration profile.
-v <i>level</i>	Displays detailed information. Enter 1 as the <i>level</i> .
-?	Shows the usage for this command.

See also

ibrix_httpshare, ibrix_httpvhost

ibrix_httpshare

Manages HTTP shares.

Description

An HTTP share provides access to data on X9000 Software file systems. HTTP clients access the shares using standard HTTP and HTTPS protocol services.

The `ibrix_httpshare` command can be used to add a share to an existing file system. You can create multiple shares having the same physical path, but with different sets of properties, and then assign users to the share that is appropriate for their usage.

NOTE: The file system must be mounted when you add the share.

Synopsis

Add an HTTP share:

```
ibrix_httpshare -a SHARENAME -c PROFILENAME -t VHOSTNAME -f FSNAME -p
dirpath [-P urlpath] [-u USERLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"browseable=true,readonly=true,..."`. To see a list of the allowable settings for the share, use `ibrix_httpshare -L`.

For the `-t` option, specify the virtual host that will handle the share.

The `-U` option allows you to specify the permissions allowed for specific users. Use commas to separate the users (for example: `USER1=r,USER2=rw,...`).

Modify an HTTP share:

```
ibrix_httpshare -m SHARENAME -c PROFILENAME -t VHOSTNAME [-f FSNAME -p
dirpath] [-P urlpath] [-u USERLIST] [-S SETTINGLIST]
```

Delete an HTTP share:

```
ibrix_httpshare -d SHARENAME -c PROFILENAME -t VHOSTNAME
```

Delete all HTTP shares associated with a file system:

```
ibrix_httpshare -d -f FSNAME
```

Display information about an HTTP share:

```
ibrix_httpshare -i SHARENAME -c PROFILENAME -t VHOSTNAME [-v level]
```

Use `-v 1` to display detailed information.

List all HTTP shares, or shares associated with a specific profile:

```
ibrix_httpshare -l [-c PROFILENAME] [-v level]
```

Use `-v 1` to display detailed information.

List HTTP shares associated with a specific file system:

```
ibrix_httpshare -l -f FSNAME [-v level]
```

Use `-v 1` to display detailed information.

List the valid settings for a share:

```
ibrix_httpshare -L
```

Options

Option	Description
-L	Lists valid HTTP share settings.
-P <i>urlpath</i>	Specifies the URL that clients will use to access the share.
-S <i>SETTINGLIST</i>	Specifies the settings that are to be applied to an HTTP share. Use a comma to separate the settings, and enclose the settings in quotation marks, such as "browseable=true,readonly=true,...". If a value is not specified for a setting, the default value will be used. To see a list of available settings for the share, use <code>ibrix_httpshare -L</code> .
-a <i>SHARENAME</i>	Adds an HTTP share.
-c <i>PROFILENAME</i>	Specifies a configuration profile.
-d	Deletes an HTTP share.
-f <i>FSNAME</i>	Specifies a file system.
-i	Displays information about HTTP shares.
-l	Lists HTTP shares.
-m <i>SHARENAME</i>	Modifies the specified HTTP share.
-p <i>dirpath</i>	Specifies the path to the share.
-t <i>vhost</i>	Specifies the name of the virtual host for the share.
-u <i>USERLIST</i>	Specifies users and their access permissions. Use commas to separate the users (for example: <code>USER1=r,USER2=rw,...</code>).
-v <i>level</i>	Displays detailed information. Enter 1 as the <i>level</i> .
-?	Shows the usage for this command.

See also

`ibrix_httpconfig`, `ibrix_httpvhost`

ibrix_httpvhost

Manages HTTP virtual hosts.

Description

An HTTP virtual host is associated with an HTTP configuration profile. The virtual host specifies the IP addresses/ports that clients will use to access the HTTP shares associated with the configuration profile. The virtual host also specifies the SSL certificate that will be used for HTTPS connections.

Synopsis

Add a virtual host:

```
ibrix_httpvhost -a VHOSTNAME -c PROFILENAME -I IP-Address:Port [-S SETTINGLIST]
```

For the -S option, use a comma to separate the settings, and enclose the settings in quotation marks, such as "sslcert=name, ...". To see a list of the allowable settings for the vhost, use `ibrix_httpvhost -L`.

For the -I option, use a semicolon to separate the IP-address:port settings and enclose the settings in quotation marks, such as "ip1:port1;ip2:port2;...". For example:

```
ibrix_httpvhost -a vhost1 -c myprofile -I "99.1.26.1:80;99.1.26.1:81"
```

Modify a virtual host

```
ibrix_httpvhost -m VHOSTNAME -c PROFILENAME -I IP-Address:Port [-S SETTINGLIST]
```

Delete a virtual host

```
ibrix_httpvhost -d VHOSTNAME -c PROFILENAME
```

Display information about a virtual host:

```
ibrix_httpvhost -i VHOSTNAME -c PROFILENAME [-v level]
```

Use -v 1 to display detailed information.

List all virtual hosts, or virtual hosts associated with a specific profile:

```
ibrix_httpvhost -l [-c PROFILENAME] [-v level]
```

Use -v 1 to display detailed information.

List the valid settings for a virtual host:

```
ibrix_httpvhost -L
```

Options

Option	Description
-I <i>IP-Address:Port</i>	Specifies the IP addresses/ports assigned to this virtual host.
-L	Lists valid HTTP virtual host settings.
-S <i>SETTINGLIST</i>	Specifies the settings that are to be applied to a virtual host. Use a comma to separate the settings, and enclose the settings in quotation marks, such as "sslcert=name, ...". To see a list of available settings for the vhost, use <code>ibrix_httpvhost -L</code> .
-a <i>VHOSTNAME</i>	Adds an HTTP virtual host.
-c <i>PROFILENAME</i>	Specifies a configuration profile.
-d	Deletes an HTTP virtual host.
-i	Displays information about HTTP virtual hosts.

Option	Description
-l	Lists HTTP virtual hosts.
-m <i>VHOSTNAME</i>	Modifies the specified HTTP virtual host.
-v <i>level</i>	Displays detailed information. Enter 1 as the <i>level</i> .
-?	Shows the usage for this command.

See also

`ibrix_httpconfig`, `ibrix_httpshare`

ibrix_license

Displays the contents of the X9000 File Serving Software license.

Description

`ibrix_license` lists the contents of the license file, which controls the operation of an X9000 Software cluster.

Synopsis

Display the license contents:

```
ibrix_license -i
```

In the output, "Segment Server" refers to file serving nodes.

Options

Option	Description
-i	Lists license contents.
-?	Shows the usage for this command.

ibrix_localgroups

Manages local group accounts for CIFS or HTTP access.

Description

If Local User authentication is configured for CIFS or HTTP shares, use the `ibrix_localgroups` command to create and manage local group accounts.

When naming local groups, you should be aware of the following:

- Group names must be unique. The new name cannot already be used by another user or group.
- The following names cannot be used: administrator, guest, root.

Synopsis

Add a local group account:

```
ibrix_localgroups -a -g GROUPNAME [-G GROUPID] [-S RID]
```

The GID and RID will be generated automatically if you do not specify values for them.

Modify a local group account:

```
ibrix_localgroups -m -g GROUPNAME [-G GROUPID] [-S RID]
```

If you are changing the GID or RID for the group, it cannot be the primary group for any local users.

Delete a local group account:

```
ibrix_localgroups -d -g GROUPNAME
```

Display information about a specific local group account:

```
ibrix_localgroups -l -g GROUPNAME
```

Display information about all local group accounts:

```
ibrix_localgroups -L
```

Options

Option	Description
-G <i>GROUPID</i>	Specifies the group ID. The range of IDs is 2000–65535.
-L	Displays information about all local groups.
-S <i>RID</i>	Specifies the RID for the local group. The RID is the last <i>n</i> digits of the SID, ranging from 2000 to 40000000.
-a	Adds a local group.
-d	Deletes a local group.
-g <i>GROUPNAME</i>	Specifies the name of the local group.
-l	Displays information about a specific local group.
-m	Modifies a local group.
-?	Shows the usage for this command.

See also

`ibrix_localusers`, `ibrix_auth`

ibrix_localusers

Manages local user accounts for CIFS or HTTP access.

Description

If Local User authentication is configured for CIFS or HTTP shares, use the `ibrix_localusers` command to create and manage local user accounts.

When naming local users, you should be aware of the following:

- User names must be unique. The new name cannot already be used by another user or group.
- The following names cannot be used: administrator, guest, root.

Synopsis

Add a local user account:

```
ibrix_localusers -a -u USERNAME -g DEFAULTGROUP -p PASSWORD [-n] [-h HOMEDIR] [-s SHELL] [-i USERINFO] [-U USERID] [-S RID] [-G GROUPLIST]
```

The UID and RID will be generated automatically if you do not specify values for them.

Modify a local user account:

```
ibrix_localusers -m -u USERNAME [-g DEFAULTGROUP] [-p PASSWORD] [-n] [-h HOMEDIR] [-s SHELL] [-i USERINFO] [-G GROUPLIST]
```

You cannot change the UID or RID for the account. If it is necessary to change a UID or RID, you will need to delete the account and then recreate it with the new UID or RID.

Delete a local user account:

```
ibrix_localusers -d -u USERNAME
```

Display information about a specific local user account:

```
ibrix_localusers -l -u USERNAME
```

Display information about all local user accounts:

```
ibrix_localusers -L
```

Options

Option	Description
<code>-G GROUPLIST</code>	Specifies additional groups for the user.
<code>-L</code>	Displays information about all local users.
<code>-S RID</code>	Specifies the RID for the local user. The RID is the last <i>n</i> digits of the SID, ranging from 2000 to 40000000.
<code>-U USERID</code>	Specifies the UID for the local user.
<code>-a</code>	Adds a local user account.
<code>-d</code>	Deletes a local user account.
<code>-g DEFAULTGROUP</code>	Specifies the user's default local group.
<code>-h HOMEDIR</code>	Specifies the user's home directory. The default is <code>/home/<username></code> .
<code>-i USERINFO</code>	Specifies user information (for example, Doe, John, tel: 555-324-1212, Manager for services group).
<code>-l</code>	Displays information about a specific local user.
<code>-m</code>	Modifies a local user account.

Option	Description
-n	Specifies that the password is in NT hashed format.
-p <i>PASSWORD</i>	Specifies the local user's password.
-s <i>SHELL</i>	Specifies the local user's shell program. The default is <code>/bin/false</code> .
-u <i>USERNAME</i>	Specifies the username for the account.
-?	Shows the usage for this command.

See also

`ibrix_localgroups`, `ibrix_auth`

ibrix_lv

Manages logical volumes.

Description

`ibrix_lv` creates logical volumes (file-system segments) from a volume group. You can specify logical volume names or allow the names to be assigned automatically. Optionally, you can specify a logical volume size (the minimum size is 1 GB). Additional memory allocation is rounded to the next higher 32 MB. If you do not specify a size, the management console evenly divides the whole volume group among the logical volumes.

Linux `lvcreate` creates a logical volume in an existing volume group. With some forms of `ibrix_lv`, you can set values for Linux `lvcreate` options.

Synopsis

Create logical volumes from the specified volume group and name them with the string *LVNAME* followed by an underscore and a numeric suffix:

```
ibrix_lv -c [-S SIZE] [-o "OPTIONS"] -n COUNT -g VGNAME -s LVNAME
```

Numeric suffixes are assigned in sequential order beginning with 1. The maximum value of the series is set by the value of *COUNT*. To specify logical volume size, include the `-S SIZE` option. For example, to create three 64-GB logical volumes named `ilv_1`, `ilv_2`, and `ilv_3` from volume group `ivg1`:

```
ibrix_lv -c -S 64000 -n 3 -g ivg1 -s ilv
```

To specify Linux `lvcreate` options, include the `-o "OPTIONS"` option. For example to set a readahead value of 60 sectors, enter `-o "-r 60"`.

Create logical volumes from the specified volume group and assign them the names in *LVLIST*:

```
ibrix_lv -c [-S SIZE] [-o "OPTIONS"] -g VGNAME -s LVLIST
```

To specify logical volume size, include the `-S SIZE` option. The following example creates two 64-GB logical volumes named `ilv1` and `ilv2` from volume group `ivg1`:

```
-c -S 64000 -g ivg1 -s ilv1,ilv2
```

To specify `lvcreate` options, include the `-o "OPTIONS"` option.

List segment information:

```
ibrix_lv -l [-s LVLIST]
```

The command reports information for all segments or the segments specified in *LVLIST*.

Delete the segments specified in *LVLIST*:

```
ibrix_lv -d -s LVLIST
```

Options

Object	Description
<code>-S <i>SIZE</i></code>	Specifies the size in MB of a logical volume. The minimum size is 1 GB. For sizes larger than 1 GB, the value is adjusted to the next higher 32 MB increment.
<code>-c</code>	Creates a logical volume.
<code>-d</code>	Deletes a logical volume.
<code>-g <i>VGNAME</i></code>	Specifies a volume group name.
<code>-l</code>	Lists logical volume information, including size in MB, name of file system, segment number, volume group name, and any options.

Object	Description
-n <i>COUNT</i>	Specifies the number of logical volumes to create. The value is an integer that specifies both the number of logical volumes to create and the maximum value of the numeric postfix for each logical volume name.
-o " <i>OPTIONS</i> "	A string of options for Linux <code>lvcreate</code> , enclosed in double quotation marks.
-s <i>LVNAME</i> or <i>LVLIST</i>	A logical volume or a list of logical volumes. Lists can be either a comma-delimited list (for example, <code>ilv1,ilv2,ilv3</code>) or a range list (for example, <code>ilv[1-3]</code>).
-?	Shows the usage for this command.

See also

`ibrix_snap`, `ibrix_vg`, `Linux lvcreate`

ibrix_migrator

Defines migration rules and controls migrator operations.

Description

Tiering enables automatic migration of files from one tier to another within the same file system. User-written rules based on file attributes (such as modification time, access time, file size, or file type) define the migration policy, determining which files are moved and when. Tiers are defined on segments, not individual files. Use `ibrix_tier` to assign segments to tiers. You can then use `ibrix_migrator` to define the tiering policy for the file system.

When configuring the tiering policy, note the following:

- The tiering policy (a set of rules) applies to individual files contained in a specific file system.
- If a file meets the criteria of a rule, it will be moved from its current tier to the rule's target tier.
- Once configured, tiering policy is executed via command or as a `cron` job, and is performed in the background.
- Recently touched files (files that have been created or changed within the past five minutes) are considered active and will not be moved.

Migrator command restrictions

When migration rules are defined, tiering operations can be started and stopped with the `ibrix_migrator` command. Only one tiering operation can be running on a file system. Tiering operations are treated as run-to-completion tasks that are not restarted on failure and cannot be paused and later resumed. However, tiering can be started if a server is in the In-Failover state.

`ibrix_migrator` cannot be run at the same time as `ibrix_rebalance`.

Synopsis

Create a rule defining data migration from a source tier to a destination tier:

```
ibrix_migrator -A -f FSNAME -r RULE -S SOURCE_TIER -D DESTINATION_TIER
```

The rule is written to the configuration database.

Delete rules from the configuration database for the specified file system:

```
ibrix_migrator -d -f FSNAME -r RULELIST
```

Use the `-f` option to limit the command to a specific file system.

List information about active migrator tasks:

```
ibrix_migrator -i [-f FSNAME]
```

Use the `-f` option to limit the command to a specific file system.

List information about migrations:

```
ibrix_migrator -l [-f FSNAME] [-r]
```

The `-r` option lists rule information.

Start a migration operation on the specified file system:

```
ibrix_migrator -s -f FSNAME [-m MININACTIVESECONDS] [-v]
```

Stop the specified migration operation:

```
ibrix_migrator -k -t TASKID [-F force]
```

Options

Option	Description
-A	Adds a data migration rule to a file system.
-D <i>DESTINATION_TIER</i>	Specifies the destination tier in a migration rule.
-F <i>force</i>	Forces migration.
-S <i>SOURCE_TIER</i>	Specifies the source tier in a migration rule.
-d	Deletes a migration rule for a file system.
-f <i>FSNAME</i>	Specifies a file system name.
-i	Lists detailed migration status information for one or more file systems, including the state of any migration tasks.
-k	Stops a migrator operation on a file system.
-l	Lists migrator information.
-m	Specifies the minimum number of inactive seconds.
-r <i>RULE</i> or <i>RULELIST</i>	A rule or list of rules.
-s	Starts a migrator operation on a file system.
-t <i>TIERNAME</i>	Identifies the tier in the current migrator operation.
-v	Provides verbose kernel logging.
-?	Shows the usage for this command.

Writing tiering rules

A tiering policy consists of one or more rules, each identifying a desired movement of files between tiers. Rules are written using the management console GUI or directly to the configuration database with the `ibrix_migrator -A` command.

This section provides definitions of rule components and examples of rules.

Operators and date/time qualifiers

The valid rules operators are: `<`, `<=`, `=`, `!=`, `>`, `>=`, plus boolean `and` and `or`.

Use the following qualifiers for fixed times and dates:

- **Time:** Enter as three pairs of colon-separated integers specifying time on a 24-hour clock. The format is `hh:mm:ss` (for example, 15:30:00).
- **Date:** Enter as `yyyy-mm-dd [hh:mm:ss]`, where time of day is optional (for example, 2008-06-04 or 2008-06-04 15:30:00). Note the space separating the date and time.

When specifying an absolute date, an absolute time, or both, the rule must use a compare type operator (`<` | `<=` | `=` | `!=` | `>` | `>=`). For example:

```
ibrix_migrator -A -f ifs2 -r "atime > '2010-09-23' " -S TIER1 -D TIER2
```

Use the following qualifiers for relative times and dates:

- **Relative time:** Enter in rules as `year` or `years`, `month` or `months`, `week` or `weeks`, `day` or `days`, `hour` or `hours`.
- **Relative date:** Use `older than` or `younger than`. The rules engine uses the time the `ibrix_migrator` command starts execution as the start time for the rule. It then computes the required time for the rule based on this start time. For example, `ctime older than 4 weeks` refers to that time period more that 4 weeks before the start time.

The following example uses a relative date:

```
ibrix_migrator -A -f ifs2 -r "atime older than 2 days " -S TIER1 -D TIER2
```

Rule keywords

Rules consist of keywords, qualifiers, and operators.

Keyword	Definition
atime	Access time, used in a rule as a fixed or relative time.
ctime	Change time, used in a rule as a fixed or relative time.
mtime	Modification time, used in a rule as a fixed or relative time.
gid	An integer corresponding to a group ID.
gname	A string corresponding to a group name. Enclose the name string in double quotes.
uid	An integer corresponding to a user ID.
uname	Value is a string corresponding to a user name, where the user is the owner of the file. Enclose the name string in double quotes.
size	In size-based rules, the threshold value for determining migration. This is an integer specified in K (KB), M (MB), G (GB), or T (TB). Do not separate the value from its unit (for example 24K).
type	The file-system entity the rule operates on. Only the file entity is supported in the current version of the product.
name	A regular expression, typically used to match file names. Enclose a regular expression in double quotes. The asterisk (*) wildcard is valid. For example: name = "*.mpg" A name cannot contain a " / " character; you cannot specify a path. Only a file name is allowed.
path	Path name that allows these wild cards: *, ?, /. For example, if the mountpoint for the file system is /mnt, path=ibfs1/mydir/* matches the entire directory subtree under /mnt/ibfs1/mydir. (A path cannot start with a /).
strict_path	Path name that rigidly conforms to UNIX shell file name expansion behavior. For example, strict_path=/mnt/ibfs1/mydir/* matches only the files that are explicitly in the mydir directory, but does not match any files in subdirectories of mydir.

Examples of migration rules

When writing a rule, identify the following:

- File system (-f)
- Source tier (-S)
- Destination tier (-D)

The rule portion of the command must be enclosed in single quotes. The format for writing a rule is:

```
ibrix_migrator -A -f FSNAME -r 'RULE' -S SOURCE_TIER -D DEST_TIER
```

Example 1: This example writes a rule based on the file's last modification time, using a relative time period. All files that were last modified more than one month ago will be moved. The rule string is always enclosed in single quotes.

```
# ibrix_migrator -A -f ifs2 -r 'mtime older than 1 month' -S T1 -D T2
```

Example 2: This example modifies this rule to limit the files being migrated to two types of graphic files. The `or` expressions are enclosed in parentheses, and the `*` wildcard character is used to match file name patterns.

```
# ibrix_migrator -A -f ifs2 -r 'mtime older than 1 month and ( name = "*.jpg"
or name = "*.gif" )' -S T1 -D T2
```

Example 3: This example uses `and` to impose three conditions on the migration. Note the use of 10M, with no space separating the integer and unit defining the size threshold.

```
# ibrix_migrator -A -f ifs2 -r 'ctime older than 1 month and type = file
and size >= 10M' -S T1 -D T2
```

Example 4: This example uses the `path` keyword. It moves files greater than or equal to 5M that are under the directory `/ifs2/tiering_test` from TIER1 to TIER2:

```
ibrix_migrator -A -f ifs2 -r "path = tiering_test and size >= 5M" -S
TIER1 -D TIER2
```

Example 5: Rules can be group- or user-based as well as time- or data-based. This example migrates files associated with two users to T2, with no consideration of time. The names are quoted strings.

```
# ibrix_migrator -A -f ifs2 -r 'type = file and ( uname = "X9000user"
or uname = "nobody" )' -S T1 -D T2
```

Example 6: Conditions can be combined with `and/or` to create precise (even overly precise) tiering rules.

```
# ibrix_migrator -A -f ifs2 -r ' (ctime older than 3 weeks and ctime younger
than 4 weeks) and type = file and ( name = "*.jpg" or name = "*.gif" )
and (size >= 10M and size <= 25M)' -S T1 -D T2
```

Example 7: Include the `-r` option in the standard list command to view rule definitions on this file system.

```
ibrix_migrator -v -r -f ifs2
```

The output lists the file-system name, the rule ID (IDs are assigned in the order in which rules are added to the configuration database), the rule definition, and the source and destination tiers. For example, the rule in Example 2 displays as:

```
ifs2 2 mtime older than 1 month and ( name = "*.jpg" or name = "*.gif" ) T1 T2
```

Example 8: Specify the rule ID (2 in this example) in the delete command to remove the rule from the database.

```
ibrix_migrator -d -f ifs2 -r 2
```

See also

`ibrix_fs`, `ibrix_tier`

ibrix_mount

Mounts a file system on file serving nodes and X9000 clients.

Description

When mounting a file system on specific file serving nodes, the first file serving node listed must own the root segment (segment 1). X9000 Software automatically mounts on the root segment first if you mount on all file serving nodes in the cluster.

Mounts are immediately executed for file serving nodes. For X9000 clients the mount intention is stored in the management console. When X9000 services start on an X9000 client, the client queries the management console for the file systems that it should mount and mounts them. If X9000 services are already running on the X9000 client when you set new mounts, you can force the client to query the management console by running `ibrix_client` or `ibrix_lvmount -a` on the X9000 client, or by rebooting it.

When mounting a file system where Export Control is enabled, you must specify that all clients have either RO or RW access to the file system. The default is RO. In addition, the root user can be restricted to read-only access on export-controlled systems in a hostgroup by adding the `root_ro` parameter to the `ibrix_mount` command.

Mountpoints must be defined with `ibrix_mountpoint` before running `ibrix_mount`.

Synopsis

Mount a file system on file serving nodes and X9000 clients:

```
ibrix_mount -f FSNAME [-o {RW|RO}] [-O MOUNTOPTIONS] [-h HOSTLIST] -m MOUNTPOINT
```

The file system is mounted at the specified mountpoint on the file serving nodes and X9000 clients specified in *HOSTLIST*. The first file serving node listed must own the root segment in the file system. To mount on all file serving nodes and X9000 clients, omit the *HOSTLIST* option. If Export Control is enabled, you must specify RW or RO file-system access.

Mount a file system on hostgroups:

```
ibrix_mount -f FSNAME [-o {RW|RO}] [root_ro] [-O MOUNTOPTIONS] -g GROUPLIST -m MOUNTPOINT
```

The file system is mounted at the specified mountpoint on the hostgroups specified in *GROUPLIST*. If Export Control is enabled, you must specify RW or RO file-system access. The optional `root_ro` parameter specifies that the root user be limited to read-only access on the systems in the hostgroup.

Mount options. The values for the `-O` option are:

`noatime`

Do not update the file access time when the file is accessed.

`nodiratime`

Do not update the directory access time when the directory is accessed.

`mountpath=xxx`

For X9000 clients only, mount on the specified subdirectory path of the file system instead of the root.

`nodquotstatfs`

Disables file system reporting based on directory tree quota limits.

If you specify multiple options, use commas to separate them:

```
ibrix_mount -O noatime,nodiratime -f ifs3 -m /ifs3
```

Export Control Examples

The following command mounts file system `ifs1` on all file serving nodes and X9000 clients at mountpoint `/ifs1` and gives the clients RW access:

```
<installdirectory>/bin/ibrix_mount -f ifs1 -o RW -m /ifs1
```

The following command mounts file system `ifs1` on hostgroup `finance` at mountpoint `/ifs1` and gives the clients RW access:

```
<installdirectory>/bin/ibrix_mount -f ifs1 -o RW -g finance -m /ifs1
```

To run the same command and grant clients in hostgroup `finance` read-write access, but limit the root user to RO access:

```
<installdirectory>/bin/ibrix_mount -f ifs1 -o RW root_ro -g finance -m /ifs1
```

Options

Option	Description
<code>-O MOUNTOPTIONS</code>	Client mount options.
<code>-f FSNAME</code>	A file system.
<code>-g GROUPLIST</code>	Specifies a list of one or more hostgroups.
<code>-h HOSTLIST</code>	Specifies a list of file serving nodes or X9000 clients.
<code>-m MOUNTPOINT</code>	Specifies a mountpoint.
<code>-o RO RW</code>	Specifies the level of client access granted on the file system. Options are Read Only or Read Write. Read Only is the default.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_lwmount`, `ibrix_mountpoint`

ibrix_mountpoint

Creates and deletes mountpoints.

Description

Mountpoints are immediately created or deleted on file serving nodes. For X9000 clients the intent to create or delete a mountpoint is stored in the management console. When X9000 Software services start on an X9000 client, the client queries the management console to obtain its mountpoints. If the services are already running on the X9000 client when you create or delete a mountpoint, you can force the client to query the management console by running `ibrix_client` or `ibrix_lwmount -a` on the client, or by rebooting it.

The command fails if the mountpoint already exists.

Synopsis

Create a mountpoint on file serving nodes and X9000 clients:

```
ibrix_mountpoint -c [-h HOSTLIST] -m MOUNTPOINT
```

Specify the appropriate file serving nodes and X9000 clients in *HOSTLIST*. To create a mountpoint on all file serving nodes and X9000 clients, omit the `-h HOSTLIST` option. The following example creates mountpoint `/ifs1` on nodes `s1.mycompany.com` and `s2.mycompany.com`:

```
<installdirectory>/bin/ibrix_mountpoint -c -h s1.mycompany.com,s2.mycompany.com  
-m /ifs1
```

Create a mountpoint on hostgroups:

```
ibrix_mountpoint -c -g GROUPLIST -m MOUNTPOINT
```

Specify the appropriate hostgroups in *GROUPLIST*. The following command creates the mountpoint `/ifs1` on the clients hostgroup:

```
<installdirectory>/bin/ibrix_mountpoint -c -g clients -m /ifs1
```

Delete a mountpoint from file serving nodes and X9000 clients:

```
ibrix_mountpoint -d [-h HOSTLIST] -m MOUNTPOINT
```

Specify the appropriate file serving nodes and X9000 clients in *HOSTLIST*. To delete a mountpoint from all file serving nodes and X9000 clients, omit the `-h HOSTLIST` option.

Delete a mountpoint from hostgroups:

```
ibrix_mountpoint -d -g GROUPLIST -m MOUNTPOINT
```

Specify the appropriate hostgroups in *GROUPLIST*.

List mountpoints:

```
ibrix_mountpoint -l [-h HOSTLIST | -g GROUPLIST] [-m MOUNTPOINTLIST]
```

To list all mountpoints on all file serving nodes, X9000 clients, and hostgroups, omit all options.

To list mountpoints for certain file serving nodes and X9000 clients, include only the *HOSTLIST* option. To list mountpoints for certain hostgroups, include only the *GROUPLIST* option.

To list all file serving nodes, X9000 clients, and hostgroups for a specific mountpoint, include only the `-m MOUNTPOINTLIST` option.

Options

Option	Description
-c	Creates a mountpoint.
-d	Deletes a mountpoint.
-g <i>GROUPLIST</i>	Specifies a list of hostgroups.
-h <i>HOSTLIST</i>	Specifies a list of file serving nodes or X9000 clients.
-l	Lists mountpoints on one or more hosts.
-m <i>MOUNTPOINTLIST</i>	Specifies a list of mountpoints.
-?	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_mount`

ibrix_ndmpconfig

Configures NDMP parameters on the cluster.

Description

The `ibrix_ndmpconfig` command sets NDMP parameters on the cluster. These parameters control communications between the DMA and the NDMP Servers hosted on the file serving nodes.

Synopsis

Set NDMP parameters:

```
ibrix_ndmpconfig -c [-d IP1,IP2,IP3,...] [-m MINPORT] [-x MAXPORT] [-n LISTENPORT] [-u USERNAME] [-p PASSWORD] [-e {0=disable,1=enable}] [-v {0-10}] [-w BYTES] [-z NUMSESSIONS]
```

Display the current value of all configurable parameters:

```
ibrix_ndmpconfig -i
```

Synchronize the current parameter values to all file serving nodes:

```
ibrix_ndmpconfig -s
```

Options

Option	Description
-c	Sets NDMP configuration parameters.
-d	Specifies IP addresses that can be used by a DMA to access the cluster.
-e	Enables or disables new NDMP sessions.
-i	Displays the current value of all NDMP configurable parameters.
-m <i>MINPORT</i>	Sets the minimum NDMP data transfer socket port. The default is 49,152, and the minimum allowed value is 1,024.
-n <i>LISTENPORT</i>	Sets the NDMP listener socket port. The default port is 10,000, and the minimum allowed value is 1,024.
-p <i>PASSWORD</i>	Sets the NDMP user password.
-s	Synchronizes the current parameter values to all file serving nodes.
-u <i>USERNAME</i>	Sets the NDMP user name.
-v	Sets the NDMP trace log level. This value should be set to 0, the default. The level should be increased only under the direction of HP Support.
-w <i>BYTES</i>	Sets the TCP window size for NDMP data transfer. This value should be changed only for performance reasons. The default value is 160,000.
-x <i>MAXPORT</i>	Sets the maximum NDMP data transfer socket port. The default value is 65,535. This is also the maximum allowed value.
-z <i>NUMSESSIONS</i>	Sets the maximum number of concurrent sessions per NDMP server. The default value is 128.
-?	Shows the usage for this command.

ibrix_ndmpsession

Manages NDMP sessions and reports session history.

Description

The `ibrix_ndmpsession` command lists the currently running NDMP sessions or only the completed sessions (successful or not). You can also use the command to cancel a specific NDMP session, to cancel all sessions on a specific file serving node, or to cancel all sessions on all nodes. Canceling a session kills all spawned sessions processes and frees their resources if necessary.

Synopsis

Display currently running NDMP sessions on all file serving nodes:

```
ibrix_ndmpsession -l
```

Display completed NDMP sessions (successful or not):

```
ibrix_ndmpsession -l -s [-t YYYY-MM-DD]
```

The `-t` option restricts the history to sessions occurring on or before the specified date.

Cancel one or more sessions on the specified file serving node:

```
ibrix_ndmpsession -c SESSION1,SESSION2,SESSION3,... -h HOST
```

Options

Option	Description
<code>-c</code>	Cancels NDMP sessions.
<code>-h HOST</code>	Specifies a file serving node.
<code>-l</code>	Displays NDMP sessions.
<code>-s</code>	Displays NDMP session history.
<code>-t YYYY-MM-DD</code>	Specifies a date for an historical report.
<code>-?</code>	Shows the usage for this command.

ibrix_nic

Configures network interfaces.

Description

`ibrix_nic` is used to:

- Set up automated failover for a user interface.
- Identify or delete network interfaces. User interfaces that are intended for X9000 client traffic must be identified. HP recommends that you also identify user interfaces that are intended for NFS traffic so that you can implement automated failover for them.
- Identify or delete standbys for a user interface.
- Set up network interface monitoring by configuring file serving nodes to monitor each other's health over the interface.
- Change the routing table entry for a network interface (X9000 Software supports only one route per interface).
- Change Linux `ifconfig` options.
- List information about network interfaces that X9000 Software is managing.
- On an existing cluster, add a second NIC to enable two paths of communication. (The preferred method is to use bonds.)
- In a cluster configuration with two communication paths, switch the roles of the primary and secondary NIC.
- Configure link state monitoring for iSCSI network interfaces.

For comprehensive information on setting up automated failover for a user interface, see the administration guide for your system.

Synopsis

Discover and add a physical or virtual network interface:

The command adds network interface *IFNAME* as the interface to the hosts in *HOSTLIST*. The command also indicates a secondary cluster interface if `-b` is specified.

```
ibrix_nic -a -n IFNAME -h HOSTLIST [-b]
```

Delete a network interface:

```
ibrix_nic -d -n IFNAME -h HOSTLIST
```

The specified network interface is deleted from the hosts in *HOSTLIST*. If you delete an interface from a file serving node, be sure to identify a new interface that the machine can use to contact the management console.

Assign a standby network interface:

```
ibrix_nic -b -H HOSTNAME1/IFNAME1,HOSTNAME2/IFNAME2
```

The command adds network interface *IFNAME2* on server *HOSTNAME2* as the standby for network interface *IFNAME1* on server *HOSTNAME1*.

Delete a standby network interface:

```
ibrix_nic -b -U HOSTNAME1/IFNAME1
```

The command removes standby interface *IFNAME1* on server *HOSTNAME1*.

Set up network interface monitoring:

```
ibrix_nic -m -h MONHOST -A DESTHOST/IFNAME
```

The command assigns server *MONHOST* to monitor server *DESTHOST* over network interface *IFNAME*.

Stop network interface monitoring:

```
ibrix_nic -m -h MONHOST -D DESTHOST/IFNAME
```

The command deletes server *MONHOST* as the monitor for server *DESTHOST* over network interface *IFNAME*.

Add a routing entry for a network interface:

```
ibrix_nic -r -n IFNAME -h HOSTNAME -A -R ROUTE
```

The command adds the specified route for network interface *IFNAME* on server *HOSTNAME* into the routing table, overwriting any existing route entry for the interface.

Remove a routing entry for a network interface:

```
ibrix_nic -r -n IFNAME -h HOSTNAME -D
```

The command removes the specified traffic route for network interface *IFNAME* on host *HOSTNAME*.

Modify ifconfig options:

```
ibrix_nic -c -n IFNAME -h HOSTNAME [-I IPADDR] [-M NETMASK] [-B  
BCASTADDR] [-T MTU]
```

The command modifies the specified options for network interface *IFNAME* on host *HOSTNAME*.

Switch the cluster interface on a server:

```
ibrix_nic -t -n IFNAME -h HOSTNAME [-b]
```

The command switches the cluster interface to interface *IFNAME* on server *HOSTNAME*. The command also indicates a secondary cluster interface if *-b* is specified.

List detailed information about network interfaces on specific hosts:

```
ibrix_nic -i [-h HOSTLIST] [-n NAME]
```

List summary information about network interfaces on specific hosts:

```
ibrix_nic -l [-h HOSTLIST]
```

Fail back to using the primary cluster interface on the specified host:

```
ibrix_nic -p -h HOSTNAME
```

Migrate a NIC:

```
ibrix_nic -s -H HOSTNAME1/IFNAME1,HOSTNAME2/IFNAME2 [-F]
```

If *-F* is specified, the migration is forced, skipping deactivation on *HOSTNAME1/IFNAME1*.

Configure link state monitoring for iSCSI network interfaces:

```
ibrix_nic -N -h HOST -A IFNAME
```

Link state monitoring is supported only for use with iSCSI storage network interfaces, such as those provided with X9300 Gateway systems.

Options

Option	Description
<i>-A DESTHOST/IFNAME</i>	Identifies the file serving node to be monitored and the interface used for monitoring.
<i>-B BCASTADDR</i>	Specifies an IP address used for network broadcast.
<i>-D DESTHOST/IFNAME</i>	If used with <i>-r</i> , disables network routing. If used with <i>-m</i> , deletes the monitor from the named host/interface.

Option	Description
-F	Forces NIC migration.
-H <i>HOSTNAME/IFNAME</i>	Identifies pairs of file serving nodes and interfaces in standby or monitoring relationships.
-I <i>IPADDR</i>	Specifies an IP address.
-M <i>NETMASK</i>	Specifies a netmask address.
-N	Enables link state monitoring. Link state monitoring is supported only for use with iSCSI storage network interfaces, such as those provided with X9300 Gateway systems.
-R <i>ROUTE</i>	Specifies an IP address used for routing network traffic over an interface.
-T <i>MTU</i>	Sets a MTU value in bytes, which defines the largest packet size a network can transmit.
-U	Unassigns a standby interface.
-a	Discovers and adds an interface on the named hosts.
-b	Indicates a standby (backup) interface. If used with -a or -t, indicates a secondary interface.
-c	Configures an interface on a named host.
-d	Deletes an interface from the named hosts.
-h <i>MONHOST</i> or <i>HOSTNAME</i> or <i>HOSTLIST</i>	Identifies a monitoring host, a host associated with an interface, or a list of hosts associated with an interface.
-i	Provides detailed information on interfaces.
-l	Lists summary information on interfaces.
-m	Sets up interface monitoring.
-n <i>IFNAME</i>	A network interface name. Can include a VIF suffix (:nnnnn).
-p	Fails back to the primary cluster interface.
-r	Configures a routing table entry.
-s	Migrates a NIC.
-t	Switches the cluster interface on a host.
-?	Shows the usage for this command.

ibrix_online_quotacheck

Checks and repairs quota files.

Description

You will need to use the `ibrix_online_quotacheck` command in the following cases:

- You turned quotas off for a user, the user continued to store data in a file system, and you now want to turn quotas back on for this user.
- You are setting up quotas for the first time for a user who has previously stored data in a file system.
- You are using directory tree quotas and you have moved a subdirectory into another parent directory outside of the directory that has the directory tree quota.

The `ibrix_online_quotacheck` command must be run from a file serving node that has the file system mounted.

Synopsis

The command has the following options:

```
ibrix_online_quotacheck [-gutnfsPvVh] [-T {target id}] [-I {quotacheck_id}] [-M {quotamonitor_host}] path1 [path2] ...
```

Following are some examples of using the command:

Run a quota check on a file system:

```
ibrix_online_quotacheck -M {quotamonitor_host} /mountpoint
```

For the `-M` option, it is best to specify the management console as the target host/IP address. This allows the quota usage for a deleted user or group to be updated. You can specify pathnames in place of `/mountpoint` if you want to target specific directories instead of an entire file system.

To use the `-M` option for a directory tree quota, include the `-P` option in the command:

```
ibrix_online_quotacheck -M {quotamonitor_host} -P /mountpoint
```

Run a quota check on a specific user having files only in the home directory:

```
ibrix_online_quotacheck -u -T userid home_dir
```

Run a quota check on a specific group:

```
ibrix_online_quotacheck -g -T groupid home_dir
```

Run a quota check on a specific directory tree:

```
ibrix_online_quotacheck -t -P diretree_path
```

Remove the directory tree quota for all files under a specific path:

```
ibrix_online_quotacheck -t -T 0 diretree_path
```

Options

Option	Description
<code>-I {quotacheck_id}</code>	The quotacheck ID given to this session.
<code>-M {quotamonitor_host}</code>	The quotamonitor host (typically the management console). The specified path must be a mountpoint.
<code>-P</code>	A directory tree quota.
<code>-T {target ID}</code>	Checks only for the specified target ID. For a directory tree quota, sets all entries to the target ID. If the target ID is 0, the directory tree quota is cleared for the specified path.
<code>-V</code>	Prints version.

Option	Description
-f	Forces a quotacheck ID override.
-g	Checks for the group quota type.
-n	Scans the device, but does not write the limits to the file system.
-s	Silent. Error are sent to <code>/var/log/messages</code> .
-t	Checks for the directory tree quota type.
-u	Checks for the user quota type.
-v	Verbose output.
-h	Shows the usage for this command.

ibrix_powersrc

Adds or deletes power sources and modifies power source attributes in the configuration database.

Description

Programmable power sources are required for automated failover and forced manual failover. The installed power sources must be identified and added to the configuration database before they can be used. X9000 Software works with APC power sources and with integrated power sources that are managed by iLO, IPMI, or OpenIPMI.

To implement automated failover or to centrally manage power sources, the management console must have LAN access to all identified power sources.

All types of power sources must be identified to the configuration database with `ibrix_powersrc -a` before they can be used. After you identify an APC power source, execute `ibrix_hostpower` to identify the slots where file serving nodes are connected. File serving nodes with integrated power sources are connected to slot 1 by default, making it unnecessary to declare a slot association for them.

If you move a file serving node to a different APC slot, unplug a file serving node from a slot, or change a power source's IP address or password, use

`ibrix_powersrc -m` to update the configuration database with the changes.

You can use `ibrix_powersrc -d` to dissociate a file serving node from a power source by either dissociating it from a slot or dissociating it from the power source itself.

Synopsis

Add an APC power source:

```
ibrix_powersrc -a -t {apc|apc_msp} -h POWERSRCNAME -n NUMSLOTS -I IPADDR
```

The power source is identified by its name, number of slots, and IP address.

Activate an IPMI or iLO power source:

```
ibrix_powersrc -a -t {openipmi|openipmi2|ilo} -h HOSTNAME -I IPADDR -u  
USERNAME -p PASSWORD [-s]
```

The command activates the power source on the file serving node identified by name and IP address. The `-u` and `-p` options are required. To have the management console skip BMC, include the `-s` option.

Activate a VM-based power source:

```
ibrix_powersrc -a -t {esx} -h HOSTNAME -I IPADDR [-u USERNAME -p  
PASSWORD] -C CONFPATH
```

The command activates the power source on the file serving node identified by name and IP address. To secure access to the power source, include the `-u USERNAME` and `-p PASSWORD` arguments.

Modify the IP address for APC power sources:

```
ibrix_powersrc -m -I IPADDR [-s] -h POWERSRCLIST
```

To have the management console skip BMC, include the `-s` option.

Modify the IP address for IPMI or iLO power sources:

```
ibrix_powersrc -m [-I IPADDR] [-u USERNAME] [-p PASSWORD] [-s] -h POWERSRCLIST
```

To modify the security settings for IPMI- or iLO-managed power sources, include the *USERNAME* and *PASSWORD* arguments. To have the management console skip BMC, include the *-s* option.

List information about power sources:

```
ibrix_powersrc -l [-h POWERSRCLIST]
```

The command lists the name, IP address, and number of slots for either all power sources or the power sources in *POWERSRCLIST*.

Delete power sources from the configuration database:

```
ibrix_powersrc -d -h POWERSRCLIST
```

Options

Option	Description
-C <i>CONFPATH</i>	Configuration path.
-I <i>IPADDR</i>	A power source IP address.
-a	Adds or activates a power source.
-d	Deletes or deactivates a power source.
-h <i>HOSTNAME</i> or <i>POWERSRCNAME</i> or <i>POWERSRCLIST</i>	A file serving node name, the name of a power source, or a list of power source names.
-l	Lists information about power sources.
-m	Modifies a power source definition.
-n <i>NUMSLOTS</i>	Specifies the number of power slots on a power source.
-p	Password used for authentication.
-s	Skips BMC configuration.
-t <i>POWERSRCTYPE</i>	Identifies the power source type.
-u <i>USERNAME</i>	Username for authentication.
-?	Shows the usage for this command.

See also

`ibrix_hostpower`

ibrix_profile

Manages kernel profiles.

Description

The `ibrix_profile` command is used to start, disable, or reset kernel profiles, or to display kernel profile statistics.

Synopsis

Start a kernel profile:

```
ibrix_profile -E [-h HOSTLIST]
```

Use the *HOSTLIST* option to limit the command to specific nodes.

Disable a kernel profile:

```
ibrix_profile -D [-h HOSTLIST]
```

Use the *HOSTLIST* option to limit the command to specific nodes.

Reset a kernel profile:

```
ibrix_profile -R [-h HOSTLIST]
```

Use the *HOSTLIST* option to limit the command to specific nodes.

Display kernel profile statistics:

```
ibrix_profile -l [-h HOSTLIST]
```

Use the *HOSTLIST* option to limit the command to specific nodes.

Options

Option	Description
-D	Disables the kernel profile.
-E	Starts the kernel profile.
-R	Resets the kernel profile.
-h <i>HOSTLIST</i>	A file serving node.
-l	Displays kernel profile statistics.
?	Shows the usage for this command.

ibrix_pv

Discovers, lists, and deletes physical volumes.

Description

This command discovers physical volumes that have been added to file serving nodes since the last time the command was executed. The command excludes partitions and volumes that belong to volume groups created for use outside of the X9000 Software cluster.

Synopsis

Discover physical volumes and add them to the configuration database:

```
ibrix_pv -a [-h HOSTLIST] [-o FILTERDEVLIST]
```

The command runs on all file serving nodes or only the nodes in *HOSTLIST*. Omit the *-o* *FILTERDEVLIST* option to discover all SCSI devices. To discover other types of devices, include the *-o* option with the applicable keywords:

- *allow_partitions*: SCSI and IDE disk partitions
- *ata*: ATA/IDE disks
- *loop*: loop devices
- *mpath*: LVM2 multipath
- *powerpath*: EMC PowerPath

The keywords for two of the multipath devices, *mpath* and *powerpath*, are mutually exclusive. In the unlikely event that one set of hosts uses Mpath and another set uses PowerPath, an additional run is necessary to discover all devices because the sets do not intersect:

```
<installdirectory>/bin/ibrix_pv -a -h HOSTLIST -o mpath  
<installdirectory>/bin/ibrix_pv -a -h HOSTLIST -o powerpath
```

List discovered physical volumes:

```
ibrix_pv -l [-h HOSTLIST]
```

The command runs on all file serving nodes or only the nodes specified in *HOSTLIST*.

List free physical volumes:

```
ibrix_pv -l [-f]
```

Delete physical volumes:

```
ibrix_pv -d -p PVLIST [-h HOSTLIST]
```

The command runs on all file serving nodes or only the nodes specified in *HOSTLIST*.

Options

Option	Description
-a	Discovers and adds physical volumes to the configuration.
-d	Deletes physical volumes from the configuration.
-f	Includes free physical volumes only.
-h <i>HOSTLIST</i>	A list of one or more hosts.
-l	Lists physical volume information on hosts.
-o <i>FILTERDEVLIST</i>	Storage device names or the keywords <i>allow_partitions</i> , <i>ata</i> , <i>loop</i> , <i>mpath</i> , or <i>powerpath</i> .

Option	Description
-p <i>PVLIST</i>	Physical volumes. Enter either a comma-delimited list of physical volumes (for example, d1 , d2 , d3) or a range list (for example, d [1-3]).
-?	Shows the usage for this command.

ibrix_rebalance

Redistributes files among segments to balance segment utilization and server workload, or evacuates all files from specific segments.

Description

Rebalancing files

The X9000 Software Rebalancer redistributes files among segments in a file system to balance segment utilization and server workload. Normally all segments are rebalanced, possibly as a cron job, but a rebalancing can be limited to specific segments. Only segments containing files can be rebalanced.

During rebalancing, X9000 Software calculates the average aggregate utilization of all of the source segments, and then moves files from sources to destinations to bring each source segment as close as possible to the calculated utilization threshold. The final per-segment usage depends on the average file size for the target file system.

If you do not specify sources or destinations, candidate segments are sorted into sources and destinations and then rebalanced as evenly as possible.

If you specify sources, all other candidate segments in the file system are tagged as destinations, and vice versa if you specify destinations. Following the general rule, X9000 Software calculates the utilization threshold from the sources and then brings the sources as close as possible to this value by evenly distributing excess files among all the destinations. If you specified sources, the end result is that only those segments are rebalanced, and the overflow is distributed among all remaining candidate segments. If you specified destinations, the end result is that all segments except the specified destinations are rebalanced, and the overflow is distributed only to the destinations.

If you specify both sources and destinations, only the specified sources are rebalanced and the overflow is distributed among only the specified destinations.

If there is not enough aggregate room in destination segments to hold the files to be moved, X9000 Software issues an error message and does not move any files. The more restricted the number of destinations, the higher the likelihood of this error.

Rebalancing runs in two modes: corrective and analytical. Corrective mode, the default, makes all rebalancing calculations and implements the results. Analytical mode, specified with the `-a` option, makes all rebalancing calculations and reports the results without actually migrating the files.

Evacuating a segment

Before retiring storage, you will need to move, or evacuate, all files from the segments on that storage to other segments in the same file system. If segments with adequate space are not available, you will need to use `ibrix_fs` to extend the file system. The evacuation operation is transparent to users or applications accessing the file system. The amount of IO activity on the segment affects the time required for the operation.

Synopsis

Rebalance segments:

```
ibrix_rebalance -r -f FSNAME [-a] [-e] [[-s SRCSEGMENTLIST] [-S  
SRCVLIST]] [[-d DESTSEGMENTLIST] [-D DESTVLIST]] [-m INACTIVESECONDS]  
[-v]
```

The command rebalances all segments or selected segments in file system *FSNAME*. The file system must be mounted before it can be rebalanced. If you do not enter source and destination segment

lists, all segments are rebalanced. If you enter source and destination segments, specified by either segment name or logical volume name, rebalancing activity is restricted to those segments.

To run the command in analytic mode, include the `-a` option. For information on interpreting this information, contact HP Support.

To evacuate a segment, include the `-e` option. When the rebalance operation completes, you can remove the storage from the cluster. The segment numbers associated with the storage are not reused.

Display a summary of rebalancing and segment evacuation tasks:

```
ibrix_rebalance -l [-f FSLIST]
```

Use *FSLIST* to limit the command to specific file systems.

List status information about rebalancing and segment evacuation tasks:

```
ibrix_rebalance -i [-f FSLIST]
```

The report lists jobs by task ID and file system and indicates whether each job is running or stopped. Jobs that are in the analysis ("Coordinator") phase are listed separately from those in the implementation ("Worker") phase. Use *FSLIST* to limit the command to specific file systems.

Stop a rebalancing or segment evacuation task:

```
ibrix_rebalance -k -t TASKID [-F]
```

The job is specified by *TASKID*. To force the job to stop, include the `-F` option. To obtain the *TASKID*, run `ibrix_rebalance -i`.

You can stop a rebalance task at any time. Stopping a task poses no risks for the file system. The management console completes any file migrations that are in process when you issue the stop command. Depending on when you stop a job, segments may contain more or fewer files than before the run began.

Options

Option	Description
<code>-D <i>DESTLVLIST</i></code>	Specifies a list of destination segments by logical volume name, such as <i>ilv1,ilv2,ilv3</i> .
<code>-F</code>	Forces a rebalancing or segment evacuation task to stop.
<code>-S <i>SRCVLVLIST</i></code>	Specifies a list of source segments by logical volume name, such as <i>ilv1,ilv2,ilv3</i> .
<code>-a</code>	Runs rebalancing in analytical mode. By default the command runs in corrective mode.
<code>-d <i>DESTSEGMENTLIST</i></code>	Specifies a list of destination segments by segment name or number, such as <i>1,2,3</i> .
<code>-e</code>	Evacuates all files from the specified segments.
<code>-f <i>FSNAME</i> or <i>FSLIST</i></code>	A file system or list of file systems.
<code>-i</code>	Lists status information for rebalancing and segment evacuation tasks.
<code>-k</code>	Stops rebalancing or segment evacuation tasks.
<code>-l</code>	Reports summary or rebalance status.
<code>-m <i>INACTIVESECONDS</i></code>	Specifies the minimum number of inactive seconds. The rebalancer will not move files that have been accessed during this period. The default is 10 minutes.
<code>-r</code>	Rebalances the file system.
<code>-s <i>SRCSEGMENTLIST</i></code>	Specifies a list of source segments by segment name or number, such as <i>1,2,3</i> .
<code>-t <i>TASKID</i></code>	Task ID.

Option	Description
-v	Verbose kernel logging.
-?	Shows the usage for this command.

ibrix_server

Performs management and control procedures for file serving nodes.

Description

Use `ibrix_server` to:

- Identify or delete a standby for a file serving node.
- Manually fail over a file serving node.
- Fail back a file serving node following either automated or manual failover.
- Turn automated failover on or off.
- Prefer or unprefer a user interface for a file serving node. By default, file serving nodes use the cluster interface.
- Power cycle a host.
- Lists host information, including hostname, name of backup server, operational state, whether automated failover is turned on, and host ID.
- Delete file serving nodes from the configuration database.

Synopsis

Identify a standby for a file serving node:

```
ibrix_server -b -h SERVERNAME1,SERVERNAME2
```

The command assigns file serving node *SERVERNAME2* as the standby for file serving node *SERVERNAME1*.

Fail over a file serving node to its standby:

```
ibrix_server -f [-p] [-M] [-N] -h SERVERNAME
```

Use this command to perform a manual failover. Depending on whether you configured server-level or segment-level standby for the file serving node, the command migrates either specific segments or all segments to the standby server.

If the file serving node is connected to a power source that has been identified to the management console, include the `-p` option to make the file serving node power down before the failover.

The `-M` option forces segment migration. The `-N` option skips the health check otherwise performed in segment migration.

Fail back a file serving node:

```
ibrix_server -f -U [-p] [-M] [-N] -h SERVERNAME
```

The `-p` option forces a power down on the backup server. The `-M` option forces segment migration. The `-N` option skips the health check otherwise performed in segment migration.

Turn automated failover on or off:

```
ibrix_server -m [-U] [-h SERVERNAME]
```

The command applies to all file serving nodes or only to file serving node *SERVERNAME*. To turn automated failover on, omit the `-U` option. To turn it off, include the `-U` option.

Delete the standby for a file serving node:

```
ibrix_server -b -U -h SERVERNAME
```

Delete one or more servers from the configuration database:

```
ibrix_server -d -h SERVERLIST
```

Prefer a network interface for server traffic:

```
ibrix_server -n -h SRCSEVER -A DESTHOST/IFNAME
```


The command assigns network interface *IFNAME* for traffic from source server *SRCSERVER* to destination host *DESTHOST*.

The network interface preference is executed immediately on the file serving nodes. When preferring a user network interface for traffic from a source host to a destination host, traffic in the reverse direction remains defaulted to the cluster interface.

Use the default cluster interface for server traffic:

```
ibrix_server -n -h SRCSERVER -D DESTHOST
```

The cluster interface will be used for network traffic from source server *SRCSERVER* to destination host *DESTHOST*.

Manipulate the power source for a file serving node:

```
ibrix_server -P {on|reset|off} -h SERVERNAME
```

You can turn the power source on or off, or you can reset it. The command does not trigger a failover event.

Recover file serving nodes:

```
ibrix_server -r [-h SERVERLIST] [-v]
```

Start, stop, or restart the CIFS, NFS, or ndmp service:

```
ibrix_server -s -t { cifs | nfs | ndmp } -c { start | stop | restart }  
[-h SERVERLIST]
```

Change the local hostname on a file serving node:

```
ibrix_server -R -h SERVERNAME -T HOSTNAME
```

Display server information:

```
ibrix_server -l [-h SERVERLIST] [-v]
```

The following information is displayed for all servers or the servers in *SERVERLIST*: hostname, standby server name, operating state, server ID, and whether automated failover is turned on. To also list IAD, kernel, and file-system version information, include the *-v* option.

Display operational and configuration information for file serving nodes:

```
ibrix_server -i [-h SERVERLIST] [-x]
```

Options

Option	Description
-A	Assigns a preferred NIC.
-D <i>DESTHOST</i>	Deletes the network interface preference established between the source host and this destination file serving node.
-M	Forces segment migration.
-N	Skips health checks.
-P <i>on off reset</i>	Executes the selected power action on the specified file serving node.
-R	Changes the system hostname.
-T <i>HOSTNAME</i>	The new hostname.
-U	Unassigns the backup for a server or stops auto-failover monitoring on the specified file serving node.
-b	Assigns a backup host.
-c	Controls services.
-d	Deletes a file serving node from the configuration database.

Option	Description
-f	Fails over to or fails back from a file serving node.
-h <i>SERVERNAME</i> or <i>SERVERLIST</i>	One or more file serving nodes.
-i	Provides detailed information about one or more file serving nodes.
-l	Lists file serving nodes and identifies any configured backups. Also specifies whether the nodes are up or down and whether auto-failover is configured.
-m	Toggles auto-failover monitoring on a file serving node.
-n	Prefers a network interface for communication between a given server pair.
-p	Forces a power-down.
-r	Recovers a file serving node.
-s	Service control.
-t { <i>cifs</i> <i>nfs</i> <i>ndmp</i> }	The name of the service.
-v	When used with -l, reports version information for the file system, IAD, and kernel.
-x	Displays additional server information.
-?	Shows the usage for this command.

See also

`ibrix_lwhost`

ibrix_sh

Manages shell tasks.

Description

The `ibrix_sh` command can be used to run the Linux shell commands `cp` and `find` on file serving nodes and save the output. You can also list shell tasks in the cluster, display information about the tasks, or delete tasks.

Synopsis

Execute a command on all file serving nodes:

```
ibrix_sh -c COMMAND -o OUTPUT_DIR
```

The results from the command are saved in the directory *OUTPUT_DIR*. The supported shell commands are `cp` and `find`. All paths associated with these shell commands, including *OUTPUT_DIR*, must be specified as absolute paths.

List shell tasks:

```
ibrix_sh -l [-f FILESYSTEM] [-h HOSTNAME]
```

The command can report all shell tasks in the cluster, or it can be restricted to a specific file system and file serving node.

Display detailed information about shell tasks:

```
ibrix_sh -i [-f FILESYSTEMS] [-h HOSTNAME]
```

The command can report all shell tasks in the cluster, or it can be restricted to a specific file system and file serving node.

Display detailed information for certain tasks:

```
ibrix_sh -i -n TASKIDS [-h HOSTNAME]
```

The tasks are specified by task ID. Use the `-h` option to limit the output to tasks running on the specified file serving node.

Stop the task with the specified task ID:

```
ibrix_sh -k -n TASKID
```

Options

Option	Description
<code>-c <i>COMMAND</i></code>	The command to be executed on the file serving nodes.
<code>-f <i>FILESYSTEM</i> or <i>FILESYSTEMS</i></code>	A file system.
<code>-h <i>HOSTNAME</i></code>	A file serving node.
<code>-i</code>	Displays detailed information about shell tasks.
<code>-k</code>	Stops the task with the specified task ID.
<code>-l</code>	Lists shell tasks.
<code>-n <i>TASKID</i> or <i>TASKIDS</i></code>	A task name.
<code>-o <i>OUTPUT_DIR</i></code>	The output directory that will receive the results of the command.
<code>-?</code>	Shows the usage for this command.

ibrix_snap

Creates or deletes file-system snapshots and cleans up after an unsuccessful snapshot session.

Description

`ibrix_snap` creates a file-system snapshot for clusters that implement supported storage systems. The snapshot replicates file-system entities and is managed exactly like a standard file system. A snapshot is active from the moment that it is created.

The first snapshot task is to define the snapshot allocation policy. To do this, run the `ibrix_snap -p` command to define the policy type and its parameters. Once set, a snapshot allocation policy needs to change only if a parameter value changes. When the allocation policy is in effect, create the snapshot with `ibrix_snap -c`.

Synopsis

Set the snapshot policy for the specified file system:

```
ibrix_snap -p -f ORIGINFSNAME [-h HOSTLIST]
```

By default, the snapshot policy is to mirror the segment ownership. For example, if the file system has segments 1–6 owned by hosts a, b, c, d, e, and f respectively, the snapshot file system will have segments 1–6 owned by hosts a, b, c, d, e, and f respectively. The `-p` option allows you to override the default policy and specify segment owners according to a round robin pattern. The segment owners are specified with the `-h HOSTLIST` option. Using the previous example, to change the ownership of the segments 1–6 in the snapshot file system to hosts a, c, e, a, c, e respectively, enter this command:

```
ibrix_snap -p -f fs2 -h a,c,e
```

Create a file-system snapshot:

```
ibrix_snap -c -n SNAPFSNAME -f ORIGINFSNAME [-M]
```

The command creates snapshot *SNAPFSNAME* from file system *ORIGINFSNAME*. The `-M` option automatically mounts the snapshot file system.

Delete the specified file-system snapshots:

```
ibrix_snap -d -f SNAPFSLIST
```

Clean up the specified invalid snapshot file systems:

```
ibrix_snap -r -f SNAPFSLIST
```

List information about file-system snapshots:

```
ibrix_snap -i [-f SNAPFSLIST]
```

The command reports information for all file systems or only for those specified in *SNAPFSLIST*.

List file-system snapshots for all file systems on all hosts:

```
ibrix_snap -l
```

Enumerate snapshot file systems:

```
ibrix_snap -l -f ORIGINFSNAME [-n STRATEGYNAME]
```

Run the automated snapshot procedure on a file system using the specified snapshot strategy:

```
ibrix_snap -A -f ORIGINFSNAME -n STRATEGYNAME
```

Options

Option	Description
-A	Runs the automated snapshot procedure.
-M	Automatically mounts a snapshot.
-c	Creates a file-system snapshot.
-d	Deletes a file-system snapshot.
-f <i>ORIGINFSNAME</i>	The file system for which a snapshot is being created.
-h <i>HOSTLIST</i>	A list of one or more hosts.
-i	Lists snapshot information.
-l	Lists file-system snapshots.
-n <i>SNAPFSNAME</i> or <i>SNAPFSLIST</i>	The name of a file-system snapshot or list of snapshots.
-p	Sets snapshot policy for a snapshot file system.
-r	Cleans up an invalid snapshot.
-?	Shows the usage for this command.

See also

`ibrix_snap_strategy`

ibrix_snap_strategy

Creates an automated snapshot strategy.

Description

If you will be taking a snapshot of a file system on a regular basis, you can create a snapshot strategy for that file system. The strategy specifies the number of snapshots to keep and the number of snapshots to mount on the system. You can also create naming templates for the snapshots and the snapshot mountpoints.

After creating a snapshot strategy, use either the `ibrix_at` command or the GUI to create a schedule for taking the snapshots. See “Creating automated snapshots using the GUI” in the *HP StorageWorks X9000 File Serving Software File System User Guide* for more information about the GUI method.

Synopsis

Create an automated snapshot strategy:

```
ibrix_snap_strategy -c -n NAME -k KEEP -m MOUNT [-T TYPE] [-N NAMESPEC]
[-M MOUNTSPEC]
```

The `-c` option creates an automated snapshot strategy. The other options are:

<code>-n NAME</code>	The name for the snapshot strategy.
<code>-T TYPE</code>	The strategy type; either <code>linear</code> (corresponds to “Regular” on the GUI) or <code>DWMGroup</code> . <code>linear</code> allows a specific number of snapshots and mountpoints on the system. <code>DWMGroup</code> allows a specific number of snapshots and mountpoints per day, week, and month.
<code>-k KEEP</code>	<p>The number of snapshots to keep per file system. For the HP StorageWorks 2000 Modular Smart Array G2, the maximum is 32 snapshots per file system. For EqualLogic arrays, the maximum is eight snapshots per file system.</p> <p>If the strategy type is <code>linear</code>, enter the number of snapshots to keep, such as <code>-k 6</code>.</p> <p>If the strategy type is <code>DWMGroup</code>, enter the number of snapshots to keep per day, week, and month. The numbers must be separated by commas, such as <code>-k 1,2,6</code>.</p>
<code>-m MOUNT</code>	<p>The number of snapshots to mount per file system. The maximum number of snapshots is seven per file system.</p> <p>If the strategy type is <code>linear</code>, enter the number of snapshots to mount, such as <code>-m 7</code>.</p> <p>If the strategy type is <code>DWMGroup</code>, enter the number of snapshots to mount per day, week, and month. The numbers must be separated by commas, such as <code>-m 2,2,3</code>. The sum of the numbers must be less than or equal to 7.</p>
<code>-N NAMESPEC</code>	Snapshot name template. The template specifies a scheme for creating unique names for the snapshots. Use the variables shown below for the template.
<code>-M MOUNTSPEC</code>	Snapshot mountpoint template. The template specifies a scheme for creating unique mountpoints to the snapshots. Use the variables shown below for the template.

Variables for snapshot name and mountpoint templates. The following variables can be used:

<code>fulldate</code>	<code>yyyy_mm_dd_HHmmz + GMT</code>
<code>shortdate</code>	<code>yyyy_mm_dd</code>
<code>type</code>	The strategy type specified with the <code>-T</code> option (either <code>linear</code> or <code>DWMGroup</code>)
<code>strategy</code>	Same as <code>type</code> but with the name of the snapshot strategy appended
<code>fsname</code>	File-system name

You can specify one of or more of these variables, enclosed in brackets (`{ }`) and separated by underscores (`_`). The template can also include text strings. Following are some sample templates.

When a snapshot is created with one of these templates, the variables will be replaced with the values shown above.

```
{fsname}_snap_{fulldate}  
snap_{shortdate}_{fsname}
```

Modify an automated snapshot strategy:

```
ibrix_snap_strategy -e -n NAME -k KEEP -m MOUNT [-N NAMESPEC] [-M  
MOUNTSPEC]
```

Delete an automated snapshot strategy:

```
ibrix_snap_strategy -d -n NAME
```

List automated snapshot strategies:

```
ibrix_snap_strategy -l [-T TYPE]
```

Display information about an automated snapshot strategy:

```
ibrix_snap_strategy -i -n NAME
```

Options

Option	Description
-M <i>MOUNTSPEC</i>	Snapshot mountpoint template.
-N <i>NAMESPEC</i>	Name of snapshot template.
-T <i>TYPE</i>	Strategy type; either <code>linear</code> or <code>DWMGroup</code> .
-c	Configures a snapshot strategy.
-d	Deletes a snapshot strategy.
-e	Edits an existing snapshot strategy.
-i	Displays detailed information about a snapshot strategy.
-k <i>KEEP</i>	The number of snapshots to keep per file system.
-l	Lists snapshot policies.
-m <i>MOUNT</i>	The number of snapshots to mount per file system.
-n <i>NAME</i>	Snapshot strategy name.
-?	Shows the usage for this command.

See also

`ibrix_at`

ibrix_snmpagent

Updates the configuration properties for the SNMP agent.

Description

The SNMP agent is created automatically. It is configured initially as an SNMPv2 agent and is off by default.

Certain SNMP parameters and the SNMP default port are the same regardless of SNMP version. The agent port is 5061 by default. *SYSCONTACT*, *SYSNAME*, and *SYSLOCATION* are optional MIB-II agent parameters that have no default values.

The *-c* and *-s* options are also common to all SNMP versions. The *-c* option turns the encryption of community names and passwords on or off. (Encryption is off by default.) The *-s* option toggles the agent on and off; it turns the agent on by starting a listener on the SNMP port, and turns it off by shutting off the listener. The default is off.

Two agent parameters are set at creation and cannot be updated:

- SysDescr (description of agent host system): management console
- SysObjectID (object id of the X9000 Software private MIB): .1.3.6.1.4.1.18997

Synopsis

Update SNMPv1 and v2:

```
ibrix_snmpagent -u -v {1|2} [-p PORT] [-r READCOMMUNITY] [-w  
WRITECOMMUNITY] [-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION] [-c  
{yes|no}] [-s {on|off}]
```

The command takes optional community names. The default *READCOMMUNITY* name assigned to the agent is *public*. No default *WRITECOMMUNITY* name is set (although the name *private* is often used).

Update SNMPv3:

```
ibrix_snmpagent -u -v 3 [-e engineId] [-p PORT] [-r READCOMMUNITY] [-w  
WRITECOMMUNITY] [-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION] [-y  
{yes|no}] [-z {yes|no}] [-c {yes|no}] [-s {on|off}]
```

The command includes an optional engine ID that overrides the default, which is the agent's host name. Because SNMPv3 employs different access methods from v1 and v2, the *-y* and *-z* options are used to determine if a v3 agent can process non-v3 read and write requests from the NMS. Processing is disabled by default. If processing will be enabled, the *-r* and *-w* options are required.

Lists SNMP agent properties:

```
ibrix_snmpagent -l
```

The information includes SNMP version in use; port in use; values for system contact, name and location; system description (always management console); SysObjectID (the X9000 Software vendor MIB, 1.3.6.1.4.1.18997); whether encryption is in effect; the names of the read and write communities; and whether the agent is turned on or off.

Options

Option	Description
-c <i>yes no</i>	Indicates whether community names or passwords are encrypted in the configuration file. Default: no encryption. When this option is set to <i>yes</i> , unencrypted community strings and passwords are encrypted and saved to a file when the agent starts. To make changes, replace the entry in the file with a new, unencrypted value. This value will automatically be encrypted when the agent restarts. If the option is set to "no", the values will not be changed even if they are in encrypted format; however, they will be decrypted by the agent.
-e <i>ENGINE_ID</i>	Specifies the agent's SNMPv3 engine ID.
-l	Lists summary information about the SNMP agent.
-n <i>SYSNAME</i>	Specifies the assigned system name (generally the fully-qualified domain name) of the this agent.
-o <i>SYSLOCATION</i>	Specifies the physical location of the MIB-II system being managed (this agent).
-p <i>PORT</i>	Specifies the port on which the agent listens for SNMP requests from the NMS. The default port is 5061.
-r <i>READCOMMUNITY</i>	In SNMPv1 and v2, identifies the community name (that is, the password) used for an SNMP <i>get</i> request. The default read community name is <i>public</i> . Used in SNMPv3 if non-v3 read requests are allowed.
-s <i>on off</i>	Indicates whether the agent is listening for SNMP requests on the SNMP port. The default value is <i>off</i> .
-t <i>SYSCONTACT</i>	Specifies the text id of a contact person responsible for the agent's host system.
-u	Updates an SNMP agent.
-v <i>1 2 3</i>	Indicates the version of the SNMP agent.
-w <i>WRITECOMMUNITY</i>	In SNMPv1 and v2, identifies the community name (that is, the password) used for an SNMP <i>set</i> request. No default write community name is set. Used in SNMPv3 if non-v3 write requests are allowed.
-y <i>yes no</i>	In SNMPv3, indicates whether a v3 agent can process v1 or v2 <i>get</i> requests. Default: no.
-z <i>yes no</i>	In SNMPv3, indicates whether a v3 agent can process v1 or v2 <i>set</i> requests. Default: no.
-?	Shows the usage for this command.

See also

`ibrix_snmptrap`

ibrix_snmpgroup

Manages SNMP groups.

Description

A group defines the access control policy on managed objects for one or more users. All users must belong to a group. Users and groups exist only in SNMPv3. Groups are assigned a security level, which enforces use of authentication and privacy, and reference views to define the set of MIB data that group members can access.

The default view, used when a read or write view is not specified, is `excludeAll`.

Synopsis

Create a group:

```
ibrix_snmpgroup -c -g GROUPNAME -s {noAuthNoPriv|authNoPriv|authPriv}  
[-r READVIEW] [-w WRITEVIEW] [-x CONTEXT_NAME] [-m {exact|prefix}]
```

The group can be assigned various SNMPv3 security parameters including the security level, read and write views, and the context name. A context is a collection of managed objects accessible by an SNMP entity. A second option, `-m`, is used to determine the context match.

Update a group:

```
ibrix_snmpgroup -u -g GROUPNAME [-s {noAuthNoPriv|authNoPriv|authPriv}]  
[-r READVIEW] [-w WRITEVIEW] [-x CONTEXT_NAME] [-m {exact|prefix}]
```

The command uses the same options as the create (`-c`) command.

Delete a group:

```
ibrix_snmpgroup -d -g GROUPNAME
```

List group settings:

```
ibrix_snmpgroup -l [-g GROUPNAME]
```

The information includes security level, context prefix, context match, read view, and write view.

Options

Option	Description
<code>-c</code>	Creates an SNMP group.
<code>-d</code>	Deletes an SNMP group.
<code>-g <i>GROUPNAME</i></code>	Identifies an SNMP group.
<code>-l</code>	Lists SNMP group properties.
<code>-m exact prefix</code>	Identifies whether a context will be matched exactly or by its prefix. Default: <code>exact</code> .
<code>-r <i>READVIEW</i></code>	Identifies the v3 read view this group can access.
<code>-s noAuthNoPriv authNoPriv authPriv</code>	Indicates the v3 security level assigned to this group. There are three options: <ul style="list-style-type: none"><code>noAuthNoPriv</code> (no user authorization or privacy encryption performed), which is the default<code>authNoPriv</code> (user authorization performed but no privacy encryption)<code>authPriv</code> (user authorization and privacy encryption performed)
<code>-u</code>	Updates an SNMP group.
<code>-w <i>WRITEVIEW</i></code>	Identifies the v3 write view this group can access.

Option	Description
-x <i>CONTEXT_NAME</i>	Identifies the v3 context name, defining the objects this group can access. Default: no name in use.
-?	Shows the usage for this command.

See also

`ibrix_snmpuser`, `ibrix_snmpview`

ibrix_snmptrap

Creates, updates, or deletes trapsinks for any SNMP version.

Description

A trapsink is the host destination where agents send traps, which are asynchronous notifications sent by the agent to the management station. A trapsink is specified by its name or IP address. X9000 Software supports multiple trapsinks; you can define any number of trapsinks of any SNMP version, but you can define only one trapsink per host regardless of version.

At a minimum, a v1 or v2 trapsink configuration requires a destination host and SNMP version. All other parameters are optional and many will pass a default value if not specified. Trapsink configuration for SNMPv3, however, is more detailed than for earlier versions, requiring a user name and additional security parameters.

The v3 security parameters are controlled as follows:

- If the authorization password is not specified, the agent assumes the trap message does not require authentication or encryption.
- If the authentication password is provided but no privacy password is specified, the agent supplies authentication information but does not encrypt the message.
- If both the authentication and privacy passwords are specified, authentication information is provided and the message is encrypted.

Synopsis

Create a v1 or v2 trapsink:

```
ibrix_snmptrap -c -h HOSTNAME -v {1|2} [-p PORT] [-m COMMUNITY] [-s {on|off}]
```

HOSTNAME identifies the host where the agent sends traps. If a port is not specified, the command defaults to port 162. If a community is not specified, the command defaults to the read-only community, public. The -s option toggles agent trap transmission on and off. The default is on.

Update v1 or v2 trapsink settings and turn SNMP trap transmission on or off:

```
ibrix_snmptrap -u -h HOSTNAME -v {1|2} [-p PORT] [-m COMMUNITY] [-s on|off]
```

Create a v3 trapsink:

```
ibrix_snmptrap -c -h HOSTNAME -v 3 [-p PORT] -n USERNAME [-j {MD5|SHA}] [-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD] [-x CONTEXT_NAME] [-s {on|off}]
```

HOSTNAME identifies the host where the agent sends traps. If a port is not specified, the command defaults to port 162. Other options define security settings. *USERNAME* is the v3 security name of the trap sender. If unspecified, the command uses MD5 as the user authentication algorithm and DES as the privacy algorithm, with no passwords. The *CONTEXT_NAME*, if used, specifies the managed objects that can be accessed by the agent. The -s option toggles agent trap transmission on and off. The default is on.

Update v3 trapsink settings and turn SNMP trap transmission on or off:

```
ibrix_snmptrap -u -h HOSTNAME -v 3 [-p PORT] -n USERNAME [-j {MD5|SHA}] [-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD] [-x CONTEXT_NAME] [-s {on|off}]
```

Security settings are the same as those described for the create command.

Delete trapsinks on specific hosts:

```
ibrix_snmptrap -d -h HOSTLIST
```

Display information about trapsinks:

```
ibrix_snmptrap -l [-h HOSTLIST]
```

For all v1 or v2 trapsinks, or for all hosts in *HOSTLIST*, the command displays the trapsink address, SNMP version, port, community and on/off state. For v3 trapsinks, the command also displays all security settings and passwords.

Options

Option	Description
-c	Creates a trapsink.
-d	Deletes a trapsink.
-h <i>HOSTNAME</i> or <i>HOSTLIST</i>	Identifies a host, by name or IP address, where a trapsink is created or updated. Trapsinks can be deleted or listed on multiple hosts by specifying a <i>HOSTLIST</i> (use commas to separate the hosts).
-j MD5 SHA	Specifies the user authentication algorithm to use. The default is MD5.
-k <i>AUTHORIZATION_PASSWORD</i>	Specifies the user authentication password. Passwords must contain at least eight characters. By default, no password is set.
-l	Lists SNMP trapsink properties.
-m <i>COMMUNITY</i>	Associates a v1 or v2 community with a trapsink. The default community is <code>public</code> .
-n <i>USERNAME</i>	Specifies a user name.
-p <i>PORT</i>	A trapsink destination port. The default port number is 162.
-s on off	State of SNMP notifications. The default value is <code>on</code> .
-u	Updates a trapsink.
-v 1 2 3	Indicates the SNMP version of the trapsink.
-x <i>CONTEXT_NAME</i>	The context name.
-y DES AES	Identifies the v3 privacy algorithm currently in use. The default is DES.
-z <i>PRIVACY_PASSWORD</i>	Specifies the privacy password for v3 encryption and decryption. Passwords must contain at least eight characters. By default, no password is set.
-?	Shows the usage for this command.

See also

`ibrix_event`, `ibrix_snmpagent`

ibrix_snmpuser

Creates, updates, or deletes SNMPv3 users.

Description

Names the user whose keys were used to authenticate and encrypt packets.

Synopsis

Create a v3 user in a group:

```
ibrix_snmpuser -c -n USERNAME -g GROUPNAME [-j {MD5|SHA}] [-k  
AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
```

Authentication and privacy settings are optional. An authentication password is required if the group has a security level of either `authNoPriv` or `authPriv`. The privacy password is required if the group has a security level of `authPriv`. If unspecified, MD5 is used as the authentication algorithm and DES as the privacy algorithm, with no passwords assigned.

Update a v3 user:

```
ibrix_snmpuser -u -n USERNAME [-g GROUPNAME] [-j {MD5|SHA}] [-k  
AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
```

Delete a v3 user:

```
ibrix_snmpuser -d -n USERNAME
```

List user information:

```
ibrix_snmpuser -l [-n USERNAME]
```

The information includes the user's name, group, authorization, and privacy algorithms, as well as the passwords for each algorithm. The command can report on all users or on a single named user.

Options

Option	Description
-c	Creates a SNMPv3 user.
-d	Deletes an SNMP user.
-g <i>GROUPNAME</i>	Specifies the SNMPv3 group to which this user is assigned.
-j MD5 SHA	Specifies the user authentication algorithm used to authorize messages sent or received on behalf of this user. The default is MD5.
-k <i>AUTHORIZATION_PASSWORD</i>	Specifies the user authentication password. Passwords must contain at least eight characters. By default, no password is set.
-l	Lists properties for one or more SNMP users.
-n <i>USERNAME</i>	Specifies an SNMP user name.
-u	Updates an SNMP user.
-y DES AES	Identifies the v3 privacy algorithm used to encrypt messages sent or received on behalf of this user. The default is DES.
-z <i>PRIVACY_PASSWORD</i>	Specifies the privacy password for v3 encryption and decryption. Passwords must contain at least eight characters. By default, no password is set.
-?	Shows the usage for this command.

See also

`ibrix_snmpgroup`, `ibrix_snmpview`

ibrix_snmpview

Adds entries to or deletes entries from the SNMP MIB view.

Description

A MIB view consists of a collection of view subtrees that can be included or excluded from the view. Each view subtree is defined by a combination of OID subtree and a bit string mask.

The management console automatically creates the `excludeAll` view, which blocks access to all OIDs. This view cannot be deleted. It is the default read and write view if a view is not specified for a group with the `ibrix_snmpgroup` command. Its catch-all OID and mask are:

```
OID = .1
Mask = .1
```

Synopsis

Add a subtree entry to a MIB view:

```
ibrix_snmpview -a -v VIEWNAME [-t {include|exclude}] -o OID_SUBTREE [-m MASK_BITS]
```

The subtree is included in the named view by default.

Delete a MIB view:

```
ibrix_snmpview -d -v VIEWNAME
```

Display information about MIB views:

```
ibrix_snmpview -l [-v VIEWNAME]
```

For all views or the specified view, the output includes all entries by name, type, OID subtree value, and mask bits value.

Options

Option	Description
-a	Adds an entry to a view.
-d	Deletes a view.
-l	Lists properties for SNMP views.
-m MASK_BITS	Specifies the mask bits paired with the OID subtree in a v3 MIB view. The mask bits identify which of the OID subtree's sub-identifiers are significant to the MIB view.
-o OID_SUBTREE	Identifies the root of the OID subtree to be included in or excluded from a MIB view.
-t include exclude	Indicates whether an OID subtree should be included in or excluded from a view. The default is to include the subtree.
-v VIEWNAME	Specifies the name of a v3 MIB view.
-?	Shows the usage for this command.

See also

`ibrix_snmpgroup`, `ibrix_snmpuser`

ibrix_stats

Lists statistics for file serving nodes.

Description

Approximately every 20 seconds, the file serving nodes report a range of operational statistics to the management console.

Synopsis

Display statistics for file serving nodes:

```
ibrix_stats -l [-s] [-c] [-m] [-i] [-n] [-f] [-h HOSTLIST]
```

The output includes all file serving nodes or only the nodes in *HOSTLIST*.

Options

Option	Description
-c	Prints CPU statistics.
-f	Prints NFS statistics.
-h <i>HOSTLIST</i>	A list of one or more file serving nodes separated by commas.
-i	Prints I/O statistics.
-l	Lists statistics.
-m	Prints memory statistics.
-n	Prints network statistics.
-s	Prints summary statistics.
-?	Shows the usage for this command.

ibrix_supportticket

Manages support tickets, which include system and X9000 software information useful for analyzing performance issues and node terminations.

Description

When a support ticket is created, the system collects the hardware and software information needed for analysis. If a node has terminated unexpectedly, output from the crash digester is also included. This information is collated into a tar file and placed in the directory `/admin/platform/diag/support/tickets/` on the active management console. You will need to send the tar file to HP Support for analysis.

The name of the tar file is `ticket_<name>.tgz`. In the filename, `<name>` is a number, for example, `ticket_0002.tgz`. If you want to delete a ticket or view details for it, you will need to supply the name assigned to the ticket.

A support ticket is created automatically if a file serving node terminates unexpectedly. You can also create a ticket manually if your cluster is experiencing issues that need to be investigated by HP Support. For more information about creating a ticket, see the administration guide for your system.

Synopsis

Create a support ticket:

```
ibrix_supportticket -c -L Label
```

The label is for your information only.

View details for a specific support ticket:

```
ibrix_supportticket -v -n Name
```

Use the `-l` option to obtain the name of the ticket.

Delete a support ticket that is no longer relevant:

```
ibrix_supportticket -d -n Name
```

The command removes all of the collected data and metadata related to the ticket from all nodes. Use the `-l` option to obtain the name of the ticket.

List all support tickets in the cluster:

```
ibrix_supportticket -l
```

Options

Option	Description
<code>-L <i>label</i></code>	Identifies a label for the support ticket.
<code>-c</code>	Creates a support ticket.
<code>-d</code>	Deletes a support ticket.
<code>-l</code>	Lists all support tickets in the cluster.
<code>-n <i>Name</i></code>	Identifies a specific support ticket.
<code>-v</code>	Displays information for the specified support ticket.
<code>-?</code>	Shows the usage for this command.

ibrix_tape

Lists or rescans tape devices

Description

The `ibrix_tape` command can be used to list tape and media changer devices available for NDMP backups. The command can also rescan for devices.

Synopsis

List all tape and media changer devices:

```
ibrix_tape -l
```

Rescan for tape and media changer devices:

```
ibrix_tape -r
```

Options

Option	Description
-l	Lists tape devices and media changer devices.
-r	Rescans for devices
?	Shows the usage for this command.

ibrix_task

Manages tasks running in the cluster.

Description

The `ibrix_task` command can be used to view information about the tasks running in the cluster. You can also pause, resume, or suspend tasks.

Synopsis

Lists all tasks in the cluster:

```
ibrix_task -l [-f FILESYSTEM] [-h HOSTNAME] [-s]
```

The command can be restricted to a specific file system and/or file serving node. When `-s` is used, the command lists scheduled operations instead.

Display detailed information about tasks:

```
ibrix_task -i [-f FILESYSTEMS] [-h HOSTNAME] [-s]
```

The `-f` option displays tasks for the specified file systems, and `-h` displays tasks running on the specified file serving node. When `-s` is used, the command lists scheduled operations instead.

Display detailed information for specific task IDs:

```
ibrix_task -i -n TASKIDS [-h HOSTNAME] [-s]
```

The `-h` option limits the output to tasks running on the specified file serving node. When `-s` is used, the command lists scheduled operations instead.

Stop the task with the specified task ID:

```
ibrix_task -k -n TASKID [-F] [-s]
```

The `-F` option forcefully stops the task. When `-s` is used, the command stops scheduled operations instead.

Pause the task with the specified task ID:

```
ibrix_task -p -n TASKID
```

Resume the task with the specified task ID:

```
ibrix_task -r -n TASKID
```

Options

Option	Description
<code>-F</code>	Forces the task with the specified task ID to stop.
<code>-f FILESYSTEM</code> or <code>FILESYSTEMS</code>	One or more file systems.
<code>-h HOSTNAME</code>	A file serving node.
<code>-i</code>	Displays detailed information about tasks.
<code>-k</code>	Stops the task with the specified task ID.
<code>-l</code>	Lists tasks running in the cluster.
<code>-n TASKID</code> or <code>TASKIDS</code>	A task name.
<code>-p</code>	Pauses the task with the specified task ID.
<code>-r</code>	Resumes the task with the specified task ID.
<code>-s</code>	Lists scheduled operations.
<code>?</code>	Shows the usage for this command.

ibrix_tier

Manages data tiers.

Description

A tier is a group of one or more segments. *Tiering* enables automatic migration of files from one tier to another within the same file system. User-written rules based on file attributes (such as modification time, access time, file size, or file type) define the migration policy, determining which files are to be moved and when. Tiers are defined on segments, not individual files.

Use of the tiering structure on any file system is optional. Tiering is off by default, and there is no “default tier.”

The `ibrix_tier` command assigns or unassigns segments to tiers, deletes segments from tiers, and lists information about tiers.

Use `ibrix_migrator` to create or delete rules defining migration policies, to start or stop tier operations, and to list information about rules and migrator operations.

Synopsis

Assign segments to a tier:

```
ibrix_tier -a -f FSNAME -t TIERNAME -S SEGLIST
```

Unassign segments from a tier:

```
ibrix_tier -u -f FSNAME [-S SEGLIST]
```

Delete a tier from a file system:

```
ibrix_tier -d -f FSNAME -t TIERNAME
```

List tiers for a file system:

```
ibrix_tier -l -f FSNAME [-t TIERNAME]
```

Include the `-t` option to see detailed information about the tiers.

List segments belonging to a tier:

```
ibrix_tier -i -f FSNAME [-t TIERNAME]
```

Options

Option	Description
<code>-S SEGLIST</code>	A list of segment numbers or a source tier, depending on the command.
<code>-a</code>	Assigns segments to a tier.
<code>-d</code>	Deletes a tier.
<code>-f FSNAME</code>	A file system.
<code>-i</code>	Lists segments belonging to a tier.
<code>-l</code>	Lists tiers.
<code>-t TIERNAME</code>	The name of a tier.
<code>-u</code>	Unassigns a segment.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_migrator`

ibrix_umount

Unmounts a file system.

Description

`ibrix_umount` unmounts a file system from file serving nodes, X9000 clients, and hostgroups. You can also locally unmount a file system from Linux X9000 clients with `ibrix_lwumount`.

File serving nodes immediately unmount a file system. For X9000 clients the intention to unmount a file system is stored on the management console. When X9000 Software services start on a X9000 client, the client queries the management console for the file systems to be unmounted and then unmounts them. If X9000 Software services are already running on a X9000 client, you can force the client to query the management console by either executing `ibrix_lwumount -a` on the client or by rebooting it.

The root segment must be the last segment unmounted. Attempting to unmount it while other segments are still mounted will fail.

If the file system has been exported using NFS, you must unexport it with `ibrix_exportfs` before you unmount it.

Synopsis

Unmount a file system:

```
ibrix_umount -f FSNAME [-h HOSTLIST | -g GROUPLIST] [-X]
```

To unmount the file system from all file serving nodes, X9000 clients, and hostgroups, specify only a file-system name. To unmount the file system from specific file serving nodes and X9000 clients, include the `HOSTLIST` option. To unmount the file system from specific hostgroups, include the `GROUPLIST` option.

Unmount a file system mounted at a specific mountpoint:

```
ibrix_umount -m MOUNTPOINT [-h HOSTLIST | -g GROUPLIST] [-X]
```

To unmount the file system from all file serving nodes, X9000 clients, and hostgroups, specify only the mountpoint. To unmount the file system from specific file serving nodes and X9000 clients, include the `HOSTLIST` option. To unmount the file system from specific hostgroups, include the `GROUPLIST` option.

Options

Option	Description
<code>-X</code>	Skips the X9000 Software IAD.
<code>-f FSNAME</code>	The file system to be unmounted.
<code>-g GROUPLIST</code>	A list of one or more hostgroups.
<code>-h HOSTLIST</code>	A list of one or more file serving nodes or X9000 clients.
<code>-m MOUNTPOINT</code>	Identifies a mountpoint.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_lwumount`, `ibrix_mount`, `ibrix_mountpoint`

ibrix_version

Reports software version numbers for file serving nodes and X9000 clients.

Description

This command reports the following information for each file serving node:

File system	Version number of the installed file system
IAD/IAS	X9000 Software IAD version number
IAD/FS	Version number of the IAD component that communicates with the file-system module
OS	The operating system
Kernel Version	Version number of the installed OS kernel
Arch	Processor architecture

The file system and IAD/FS output fields should show matching version numbers unless you have installed special releases or patches. If the output reports mismatched version numbers and you do not know of any reason for the mismatch, contact HP technical support. A mismatch might affect the operation of your cluster.

Synopsis

List version information for file serving nodes:

```
ibrix_version -l [-h HOSTLIST]
```

Use the `-h HOSTLIST` option to limit the information to specific nodes.

List X9000 Software version information for all file serving nodes:

```
ibrix_version -l -S
```

List version information for all X9000 clients:

```
ibrix_version -l -C
```

The information includes the client hostname, file system, IAD/IAS, IAD/FS, and operating system.

Options

Option	Description
<code>-h <i>HOSTLIST</i></code>	A list of one or more hosts.
<code>-l</code>	Lists version information.
<code>-C</code>	All X9000 clients.
<code>-S</code>	All file serving nodes.
<code>-?</code>	Shows the usage for this command.

ibrix_vg

Creates, extends, reduces, or deletes volume groups.

Description

Use `ibrix_vg` to perform these tasks:

- Create volume groups from discovered physical volumes. This is the second step in the manual method of creating a file system.
- Extend or reducing a volume group by adding or removing physical volumes.
- Delete unneeded volume groups. Before deleting a volume group, first delete the logical volumes (segments) created from it.
- List information about volume groups.

Synopsis

Create a volume group:

```
ibrix_vg -c -g VGNAME -p PVLIST
```

The command creates volume group *VGNAME* from the physical volumes in *PVLIST*. For example, the following command creates volume group `ivg1` from physical volumes `d1`, `d2`, and `d3`:

```
<installdirectory>/bin/ibrix_vg -c -g ivg1 -p d[1-3]
```

Extend a volume group:

```
ibrix_vg -e -g VGNAME -p PVLIST
```

The command extends volume group *VGNAME* with the physical volumes in *PVLIST*.

Remove physical volumes from a volume group:

```
ibrix_vg -r -g VGNAME -p PVLIST
```

The command reduces volume group *VGNAME* by removing the physical volumes in *PVLIST*.

List volume groups for file serving nodes:

```
ibrix_vg -l [-h HOSTLIST]
```

Use the *HOSTLIST* option to limit the output to specific nodes.

List size and utilization information for volume groups:

```
ibrix_vg -l [-g VGLIST]
```

Use the *VGLIST* option to limit the output to specific volume groups.

List volume group information:

```
ibrix_vg -i [-g VGLIST]
```

For all volume groups or the groups specified in *VGLIST*, the command lists size and utilization information, the physical volumes used to create the groups, the logical volumes (segments) built on the groups, and the file serving nodes that can access the groups.

Delete volume groups:

```
ibrix_vg -d -g VGLIST
```

Back up the `vgconfig` maps on the management console:

```
ibrix_vg -b -g VGLIST
```


Options

Option	Description
-b	Backs up the <code>vgconfig</code> maps on the management console.
-c	Creates a volume group.
-d	Deletes a volume group.
-e	Extends a volume group.
-g <i>VGNAME</i> or <i>VGLIST</i>	A volume group or list of volume groups. Takes either a list of volume groups separated by commas (for example, <code>ivg1, ivg2, ivg3</code>) or a range of volume groups (for example, <code>ivg[1-3]</code>).
-h <i>HOSTLIST</i>	A list of one or more host names.
-i	Displays information about volume groups and the file serving nodes accessing them.
-l	Lists volume groups.
-p <i>PVLIST</i>	A list of physical volumes. Takes either a list of physical volumes separated by commas (for example, <code>d1, d2, d3</code>) or a range of physical volumes (for example, <code>d[1-3]</code>).
-r	Removes one or more physical volumes from a volume group.
-?	Shows the usage for this command.

ibrix_vs

Identifies storage resources, discovers them, and maps them to physical volumes.

Description

`ibrix_vs` identifies vendor storage resources, discovers LUNs identified in those resources, and maps them to logical physical volumes in the configuration database. Vendor storage resources must be identified and discovered before they can be used to create file-system snapshots.

Synopsis

Register a storage array:

```
ibrix_vs -r -n STORAGENAME -t {eva | exds | lefthand | msa | eqlogic}  
-I IP(s) -U USERNAME [-P PASSWORD] [-s suffix]
```

Remove the registration for the specified storage:

```
ibrix_vs -d -n STORAGENAME
```

Discover LUNs on the specified storage:

```
ibrix_vs -a [-n STORAGENAME]
```

List all storage that has been registered:

```
ibrix_vs -l
```

List detailed information for the specified storage:

```
ibrix_vs -i [-n STORAGENAMELIST]
```

Set the proxy IP address for the specified storage:

```
ibrix_vs -c -I PROXYIP -n STORAGENAME
```

Options

Option	Description
-I	IP address.
-P <i>PASSWORD</i>	Password.
-U <i>USERNAME</i>	User name.
-a	Discovers LUNs.
-c	Sets the proxy IP address.
-d	Removes the storage registration.
-i	Lists detailed storage information.
-l	Lists storage that has been registered.
-n <i>STORAGENAME</i> or <i>STORAGENAMELIST</i>	Storage system name or a list of system names.
-r	Registers storage.
-s <i>suffix</i>	Suffix.
-t <i>TYPE</i>	Storage system type (eva, exds, lefthand, msa, or eqlogic).
-?	Shows the usage for this command.

3 Commands for Linux X9000 clients and file serving nodes

The commands in this chapter are run directly on Linux X9000 clients and/or file serving nodes. The commands can be executed by any user. In a typical installation, the commands are run from the working directory `/usr/local/ibrix/bin`.

ibrix_df

Reports disk space utilization information for a file system. This command can be run on file serving nodes or X9000 clients.

Description

`ibrix_df` reports disk space utilization for the specified file system.

Synopsis

Report disk space utilization for a file system:

```
ibrix_df -f [fmname] :fsname
```

The file-system name is always required. If this X9000 client is registered with multiple management consoles, the management console name must also be entered. Specify the management console and file system as a colon-separated pair: `fmname:fsname`. For example:

```
ibrix_df -f FM01:ifs1 -F
```

Options

Option	Description
<code>-f <i>fsname</i> or <i>fmname:fsname</i></code>	Specifies either a file-system name or a paired management console name and file-system name.
<code>-?</code>	Shows the usage for this command.

Output

The `ibrix_df` report includes the following fields:

Field	Definition
Name	File-system name.
CAPACITY	Number of blocks in the file system.
FREE	Number of unused blocks of storage.
AVAIL	Number of blocks available for user files.
USED PERCENT	Percentage of total storage occupied by user files.
FILES	Number of files that can be created in the file system.
FFREE	Number of unused file inodes in the file system.

ibrix_lwhost

Tunes Linux X9000 clients. This command can be run on X9000 clients only.

Description

Use `ibrix_lwhost` to set Linux X9000 client parameters:

- Server settings
- Preferred network interface for communication with a file serving node
- Communications protocol
- Number of server threads

Changes take effect immediately and are persistent. This command also initiates a management console query that allows the X9000 client to pick up new host tunings.

Use `ibrix_lwhost --list` to list local server settings that have been changed. The output does not list default settings that are still in effect.

CAUTION: Changing tuning parameters will alter file-system performance. Contact HP technical support before changing any tuning parameters.

Synopsis

Tune server settings:

```
ibrix_lwhost --tune -p param -v value
```

Contact HP technical support for information about the values to use for *param* and *value*.

Prefer a network interface for communicating with a file serving node:

```
ibrix_lwhost --nic -h host -n server_ifname
```

Specify a communications protocol:

```
ibrix_lwhost --protocol -p {tcp|udp}
```

The default protocol is UDP.

Set the number of server threads:

```
ibrix_lwhost --threads -v value
```

Initiate a management console query to allow the X9000 client to pick up new host tunings:

```
ibrix_lwhost --a
```

For example, you could use this option if file allocation policies have been changed since the X9000 client started and you want the X9000 client to acquire the new settings immediately.

List X9000 client parameters that have been changed from their defaults:

```
ibrix_lwhost --list
```

Options

Option	Description
--a	Initiates a management console query.
-h <i>host</i>	Specifies the name of the file serving node to which the client will connect.
-n <i>server_ifname</i>	Name of a preferred network interface.
-p <i>param</i>	A tuning or protocol parameter.

Option	Description
<code>-v value</code>	Specifies a tuning parameter value or, in the case of threads, the number of server threads to set.
<code>--list</code>	Lists parameter settings that have been changed from the defaults.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_host_tune`

ibrix_lwmount

Mounts a file system on a Linux X9000 client. This command can be run on X9000 clients only.

Description

A file system must be mounted on the file serving node that owns the root segment (that is, segment 1) before it can be mounted on an X9000 client.

If the Export Control feature is not enabled, use this command to mount a file system on an X9000 client. If the Export Control feature is enabled, you cannot use this command on a X9000 client until you have first mounted the file system using `ibrix_mount`. (For more information about Export Control, see the *HP StorageWorks X9000 File Serving Software File System User Guide*.)

`ibrix_lwmount` also initiates the management console query that the X9000 client uses to obtain new mountpoint and file allocation policy settings.

Synopsis

Query the management console for new file-system mountpoints and allocation policy settings:

```
ibrix_lwmount -a
```

Mount a file system at the specified mountpoint:

```
ibrix_lwmount -f [fmname:]fsname -m mountpoint [-o options]
```

To specify Linux mount options, include the `-o options` option. The file-system name is always required. A management console name is required only if this X9000 client is registered with multiple management consoles. Specify the management console and file system as a colon-separated pair: *fmname:fsname*. For example:

```
ibrix_lwmount -f FM01:ifs1 -m /ifs1
```

Options

Option	Description
<code>-a</code>	Initiates a management console query.
<code>-f <i>fsname</i> or <i>fmname:fsname</i></code>	Specifies a file-system name or a paired management console name and file-system name. A management console name is required only when a X9000 client is registered with multiple management consoles.
<code>-m <i>mountpoint</i></code>	A specified mountpoint.
<code>-o <i>options</i></code>	One or more Linux mount options.
<code>-?</code>	Shows the usage for this command.

See also

`ibrix_hostgroup`, `ibrix_lwumount`, `ibrix_mount`

ibrix_lwumount

Detaches a file system from a Linux X9000 client. This command can be run on X9000 clients only.

Synopsis

Detach a file system from a client:

```
ibrix_lwumount -f [fmname:] fsname
```

The file-system name is always required. A management console name is required only if this X9000 client is registered with multiple management consoles. Specify the management console and file system as a pair separated by a colon: *fmname:fsname*.

Detach the file system mounted at the specified mountpoint:

```
ibrix_lwumount -m mountpoint
```

Options

Option	Description
<i>-f fsname</i> or <i>fmname:fsname</i>	A file-system name, or a paired management console name and file-system name. A management console name is required only when the X9000 client is registered with multiple management consoles.
<i>-m mountpoint</i>	A specified mountpoint.
<i>-?</i>	Shows the usage for this command.

See also

`ibrix_lwumount`

ibrix_top

Collects and presents live views of file system, host, and segment activity. The command can be run on either file serving nodes or X9000 clients.

Description

`ibrix_top` collects performance metrics for file systems, hosts and segments; and displays the results to the standard output in a form similar to that of the UNIX `top(1)` command. When the command is first run, some fields are empty or read zero (0). These fields are populated when the command is first refreshed.

The command collects data using snapshots. Specifying `-s` saves the snapshots when the session ends. By default, `ibrix_top` creates a temporary directory and writes the snapshots to it. The directory and snapshots are deleted when the `ibrix_top` session ends. Specifying `-d` and a directory name creates the snapshots in that directory, and that directory is not deleted when the session ends.

If using the temporary directory, the directory is saved with the snapshot files, and the directory name is posted so the snapshots can be accessed later.

The following table shows the statistics displayed in each view. Any data reported on a per-second basis is derived by dividing the cumulative total by the number of seconds in the collection interval. The default interval is five seconds. It can be changed by specifying a new value with `-t`. The larger the cluster, the longer it takes to collect all of the data and the longer the interval should be.

View	Statistics Displayed
Default/File system	Per file system: number of segments, percentage of blocks free, percentage of files free (where files equals inodes), KB/sec read, KB/sec written, number of new files created/sec, number of files deleted/sec, and number of file-system errors/sec.
Default/Tasks	Indicates the four busiest X9000 software tasks; reports on user and system threads.
Default/Network device	Per NIC: KB/sec sent, KB/sec received, bytes/packet sent, bytes/packet received.
Host	Per host: IP address, KB/sec read, KB/sec written, number of new files created/sec, number of files deleted/sec, and number of file-system errors/sec. All metrics also reported as an aggregate over the interval.
Segment	Per local segment: volume group, logical volume, KB/sec read, KB/sec written, number of new files created/sec, number of files deleted/sec, and number of file-system errors/sec.
Help	Identifies the keys to press to change the view and to scroll within a view.

Synopsis

Collect and present live views of file system, host, and segment activity:

```
ibrix_top [-d DIRECTORY] [-s] [-r {host|segment|help}] [-t SEC]
```

Use the `-d` option to specify a directory for the collected data. Use the `-s` option to save the data collection snapshots when the `ibrix_top` session ends.

To change the initial view, specify `-r` and a view name. If you change from the default view and later want to go back to it, enter `-r default`.

To set the collection interval, specify `-t` and a number of seconds.

Options

Option	Description
-d <i>DIRECTORY</i>	Specifies the directory where data captures are written. HP recommends that you name this directory with an absolute path.
-r default host segment help	Specifies the initial view for an <code>ibrix_top</code> session.
-s	Creates persistent data snapshots.
-t <i>SEC</i>	Sets the number of seconds in the refresh interval between data collections. The default is 5 seconds.
-?	Shows the usage for this command.

4 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Related information

The following documents provide related information:

- *HP StorageWorks X9000 File Serving Software Release Notes*
- *HP StorageWorks X9000 File Serving Software File System User Guide*
- *HP StorageWorks X9300 Network Storage Gateway Administrator Guide*
- *HP StorageWorks X9320 Network Storage System Administrator Guide*
- *HP StorageWorks X9720 Network Storage System Administrator Guide*
- *HP StorageWorks X9000 File Serving Software Installation Guide*

Related documents are available on the Manuals page at <http://www.hp.com/support/manuals>.

On the Manuals page, select **storage** > **NAS Systems** > **NAS/Storage Servers** > **HP StorageWorks X9000 Network Storage Systems**.

HP websites

For additional information, see the following HP websites:

- <http://www.hp.com>
- <http://www.hp.com/go/x9000>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support/manuals>

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

<http://www.hp.com/go/e-updates>

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

Glossary

ACE	access control entry.
ACL	access control list.
ADS	Active Directory Service.
ALB	Advanced load balancing.
BMC	Baseboard Management Configuration.
CIFS	Common Internet File System. The protocol used in Windows environments for shared folders.
CLI	Command-line interface. An interface comprised of various commands which are used to control operating system responses.
CSR	Customer self repair.
DAS	Direct attach storage. A dedicated storage device that connects directly to one or more servers.
DNS	Domain name system.
FTP	File Transfer Protocol.
GSI	Global service indicator.
HA	High availability.
HBA	Host bus adapter.
HCA	Host channel adapter.
HDD	Hard disk drive.
IAD	HP X9000 Software Administrative Daemon.
iLO	Integrated Lights-Out.
IML	Initial microcode load.
IOPS	I/Os per second.
IPMI	Intelligent Platform Management Interface.
JBOD	Just a bunch of disks.
KVM	Keyboard, video, and mouse.
LUN	Logical unit number. A LUN results from mapping a logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN.
MTU	Maximum Transmission Unit.
NAS	Network attached storage.
NFS	Network file system. The protocol used in most UNIX environments to share folders or mounts.
NIC	Network interface card. A device that handles communication between a device and other devices on a network.
NTP	Network Time Protocol. A protocol that enables the storage system's time and date to be obtained from a network-attached server, keeping multiple hosts and storage devices synchronized.
OA	HP Onboard Administrator.
OFED	OpenFabrics Enterprise Distribution.
OSD	On-screen display.
OU	Active Directory Organizational Units.
RO	Read-only access.
RPC	Remote Procedure Call.
RW	Read-write access.
SAN	Storage area network. A network of storage devices available to one or more servers.
SAS	Serial Attached SCSI.

SELinux	Security-Enhanced Linux.
SFU	Microsoft Services for UNIX.
SID	Secondary controller identifier number.
SNMP	Simple Network Management Protocol.
TCP/IP	Transmission Control Protocol/Internet Protocol.
UDP	User Datagram Protocol.
UFM	Voltaire's Unified Fabric Manager client software.
UID	Unit identification.
USM	SNMP User Security Model.
VACM	SNMP View Access Control Model.
VC	HP Virtual Connect.
VIF	Virtual interface.
WINS	Windows Internet Naming Service.
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.
WWNN	World wide node name. A globally unique 64-bit identifier assigned to each Fibre Channel node process.
WWPN	World wide port name. A unique 64-bit address used in a FC storage network to identify each device in a FC network.

Index

A

- Active Directory authentication, 14
- Active Directory, configure on Windows X9000 clients, 11
- agile management console, 43
- Autoconnect table, NFS clients, 16
- automated snapshots, 118
 - create schedule, 13

C

- certificates, SSL, 17
- CIFS, add or delete shares, 22
- CIFS, global settings, 24
- CIFS, manage authentication, 14
- CLI
 - command syntax, 10
 - working directory, 10
- cluster
 - configure event notifications, 35
 - event messages, 35
 - set configuration parameters, 29
 - tasks, manage, 132
 - view configuration, 29
- cluster virtual interface, 43
- command syntax, 10
- contacting HP, 146

D

- data tiering
 - manage, 89, 133
 - tiering rules, 90
- directory tree quotas, 53
- disk space utilization, 139
- document
 - related documentation, 146

E

- email event notification, 35
- events, configure notifications, 35

F

- file allocation, 54
 - allocation policies, 55
- file serving nodes
 - health checks, 68
 - high availability, 64
 - manage, 112
 - physical devices, discover, 107
 - power source, failover, 76
 - statistics, 129
 - tuning parameters, 71
 - version numbers, 135
- file systems
 - activity snapshots, 144
 - add or delete CIFS shares, 22

- analyze with ibrix_fsck, 58
 - automated snapshots, 118
 - control access for X9000 clients, 41
 - create, 48
 - data tiering, 89, 133
 - delete, 50
 - disable 32-bit mode, 51
 - disk space utilization, 139
 - enable or disable Export Control, 51
 - enable or disable quotas, 50
 - export replication target, 38
 - export via NFS, 41
 - extend, 49
 - file allocation, 54
 - identify standby node, 51
 - manage, 48
 - manage quotas, 32
 - migrate segments, 50
 - mount, 93
 - mount on X9000 clients, 142
 - mountpoints, 95
 - rebalance files, 109
 - resynchronize configuration on hosts, 31
 - snapshots, 116
 - unexport replication target, 38
 - unmount, 134
 - unmount on X9000 clients, 143
 - view information, 49
- FTP service
- profiles, 60
 - shares, 62
 - SSL certificates, 17

H

- HBAs
 - manage, 66
 - port monitoring, 66
- health check sets, 70
- help
 - obtaining, 146
- high availability, configure, 64
- hostgroups, manage, 74
- hosts, activity snapshots, 144
- HP
 - technical support, 146
- HP websites, 146
- HTTP service
 - configuration profile, 77
 - shares, 79
 - SSL certificates, 17
 - virtual host, 81
- HTTP, manage authentication, 14

K

- kernel profiling, 106

L

- license file, display, [83](#)
- Linux X9000 client, register, [26](#)
- Local User authentication
 - configure, [14](#)
 - local groups, [84](#)
 - local users, [85](#)
- logical volumes, manage, [87](#)

M

- management console, agile
 - control, [43](#)
 - set mode, [43](#)
- management console, tuning parameters, [45](#)
- media changer devices for NDMP backups, [131](#)
- migration, files, [89](#)
- mountpoints, create or delete, [95](#)

N

- NDMP backups
 - parameters, set, [97](#)
 - sessions, [98](#)
 - tape and media changer devices, [131](#)
- network interfaces, configure, [99](#)
- NFS clients
 - manage Autoconnect table, [16](#)
 - mount file systems automatically, [16](#)
- NFS, export file systems, [41](#)

P

- physical volumes, discover or delete, [107](#)
- power sources
 - failover, [76](#)
 - manage in configuration database, [104](#)

Q

- quotas
 - check or repair quota files, [102](#)
 - directory tree, [53](#)
 - manage, [32](#)

R

- rebalancing, files, [109](#)
- related documentation, [146](#)
- remote replication
 - export target file system, [38](#)
 - manage host and NIC preferences, [39](#)
 - monitor activity, [19](#)
 - register a remote cluster, [28](#)
 - start or stop, [19](#)
 - unexport target file system, [38](#)

S

- scheduling commands, [13](#)
- segments
 - activity snapshots, [144](#)
 - rebalance files, [109](#)
- servers, manage, [112](#)
- shell tasks, [115](#)

SMB signing, CIFS, [24](#)

- snapshots
 - allocation policy, [116](#)
 - automated, [118](#)
 - manage, [116](#)

SNMP

- agent, [120](#)
- groups, [122](#)
- MIB view, [128](#)
- SNMPv3 users, [126](#)
- trapsinks, [124](#)

SNMP event notification, [35](#)

SSL certificates, [17](#)

statistics, file serving nodes, [129](#)

storage resources, manage, [138](#)

Subscriber's Choice, HP, [146](#)

support tickets, [130](#)

T

- tape devices for NDMP backups, [131](#)
- tasks, manage, [132](#)
- technical support
 - HP, [146](#)
- tuning parameters
 - file serving nodes, [71](#)
 - X9000 clients, [71](#)
- tuning parameters, management console, [45](#)

U

user network interface, prefer or unprefer, [26](#)

V

volume groups, manage, [136](#)

W

- websites
 - HP Subscriber's Choice for Business, [146](#)

X

- X9000 clients
 - configure Active Directory, [11](#)
 - delete, [26](#)
 - health checks, [68](#)
 - hostgroups, [74](#)
 - prefer or unprefer user interface, [26](#)
 - tuning parameters, [71](#)
 - view information, [26](#)
- X9000 clients, Linux
 - mount file system, [142](#)
 - register on management console, [26](#)
 - tune, [140](#)
 - unmount file system, [143](#)
- X9000 Software, version numbers, [135](#)